Audun Jøsang. Online Reputation Systems for the Health Sector. **Electronic Journal of Health Informatics** 2008, vol 3(1). Special issue on Privacy and Security.

Online Reputation Systems for the Health Sector[*]

Audun Jøsang
Faculty of Information Technology,
Queensland University of Technology,
Brisbane, Australia

Abstract

People who are seeking medical advice and care often find it difficult to obtain reliable information about the quality and competence of health service providers. While transparent quality evaluation of products and services is commonplace in most commercial services, public access to information about the quality of health services is usually very restricted.

Online reputation and rating systems represent an emerging trend in decision support for service consumers. Reputation systems are based on collecting information about other parties in order to derive measures of their trustworthiness or reliability on various aspects. More specifically these systems use the Internet for the collection of ratings and for dissemination of derived reputation scores. Online rating systems applied to the health sector are already emerging.

This article describes robust principles for implementing online reputation systems in the health sector. In order to prevent uncontrolled ratings, our method ensures that only genuine consumers of a specific service can rate that service. The advantage of using online reputation systems in the health sector is that it can assist consumers when deciding which health services to use, and that it gives an incentive for high quality health services among health service providers.

Keywords: Reputation Systems, Trust, Reliability, Quality of Service, Health, Medical

1. Introduction

Online reputation systems represent an important type of trust management mechanisms. Such systems, which are attracting strong interest from industry and the academic research community, are increasingly being integrated with online services and applications. The advantage of these systems is that they allow service consumers to share their experience, and use this information to make better decisions about which services can be safely accessed without risking damages from poor quality or even deceptive services.

There is a great potential for online trust and reputation systems to also be applied to physical world services, such as health services. The rationale is the same, namely to use the effectiveness of the Internet to allow consumers of health services to share their experience.

The health industry is controlled by various regulations and procedures such as mandatory approval of drugs by government bodies before they can be released to the market, and procedures for handing complaints about medical negligence. The problem with most of these well intentioned schemes is the lack of public transparency. For example, in most countries and states around the world including Australia, the public does not have access to information about complaints and sanctions against medical practitioners. In addition, information about the (in)effectiveness, or about negative side effects of drugs as experienced by patients is hard to obtain, even for medical professionals. Online reputation systems have the potential to fill this gap.

Online rating systems for medical doctors are already emerging. Examples are Healthcare Reviews[1], Vimo[2] and RateMDs.com[3] which are mostly US-centric, but which can also cover other regions such as Europe, Canada, Australia and New Zealand. The websites typically allow people to rate medical doctors or other health related services on various aspects. People can also provide written comments. The ratings and comments can be anonymous. These sites only represent an addition to the already well established trend of websites that allow people to rate members of specific professions such as school teachers, university professors and lawyers. Many members of these professions feel threatened by this trend, but have few means of resisting it. The US Federal Communications Decency Act (1996) protects Internet service providers from lawsuits over materials posted by third parties. The same principle is expressed in the Australian Broadcasting Services Amendment (Online Services) Act (1999). As long as the rating websites make a reasonable effort to prevent racist or obscene comments to remain online they operate within the law.

Despite the legality of websites for rating health services, they have been criticised by members of the medical profession. For example, the Australian Medical Association Queensland (AMAQ) called for the RateMD website to be closed because in their opinion it is unacceptable that patients can anonymously rate doctors without having to defend their criticism [1]. The AMAQ's view is that complaints about doctors should be addressed to the traditional control bodies such as the Health Quality Complaints Commission.

Reputation systems represent mechanisms for trust management [2], and our view is that rating of health services can be a valuable method for managing trust in the health industry. Trust can be described as a positive state of mind caused by the perception that the risk resulting from collaborating with the trusted party is acceptable. Reputation systems enable relying parties to determine the trustworthiness of remote parties through computer mediated communication and collaboration. These mechanisms should also allow trustworthy entities to be recognised as such. The idea is that reputation systems will enable highly trustworthy entities to attract collaboration from others, and discourage low quality and fraudulent players from participating in the community, or alternatively encourage them to behave in a more trustworthy manner. Applied to the health sector, reputation systems will make high quality health service providers more attractive to patients and consumers, while it provides a strong incentive to improve quality for health service providers that perform less well.

This paper outlines the potential of using reputation systems in the health industry, and focuses on principles and architectures for implementing reputation systems in that environment. Our description includes mechanisms for preventing abuse of such systems.

2. Trust and Reputation

Manifestations of trust are easy to recognise because we experience and rely on it every day. At the same time trust is quite challenging to define because the term is being used with a variety of meanings [3].

Two types of trust which we will call *reliability trust* and *decision trust* respectively are commonly assumed in the literature.

As the name suggest, reliability trust can be interpreted as the reliability of something or somebody, and the definition by Gambetta [4] provides an example of how this can be formulated:

Reliability Trust:
Trust is the subjective probability by which an individual, A, expects that another individual, B, performs a given action on which its welfare depends.

However, trust can be more complex than Gambetta's definition indicates. For example, Falcone & Castelfranchi [5] recognise that having high (reliability) trust in a person in general is not necessarily enough to decide to enter into a situation of dependence on that person. In [5] they write: *"For example it is possible that the value of the damage per se (in case of failure) is too high to choose a given decision branch, and this independently either from the probability of the failure (even if it is very low) or from the possible payoff (even if it is very high). In other words, that danger might seem to the agent an intolerable risk."*

It may seem as a paradox that a reliable party might still be distrusted. However, the decision to trust something or somebody is not only depended on the perceived reliability of the trusted party, but also on various contextual factors. A more general definition is needed in order to have a notion of trust that also covers such situations.

Decision Trust:
Trust is the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible.

A formal model of decision trust that explicitly includes the risk of a particular transaction and the relying party's individual risk attitude in addition to the reliability of the trusted party has been proposed in [6]. Such a measure must necessarily be more complex that measures of reliability trust because of its dependence on gains, investment values and possibly other context-dependent parameters. We will not use the notion of decision trust in this study.

A question which naturally arises when interpreting trust as reliability is whether it would be better to simply use the concept of reliability because trust does not add any new information. In fact the concept of trust is in many respects redundant for modelling economic interactions [7] because well known notions such as reliability, utility and risk are adequate and sufficient for that purpose. We agree with that view, and believe that trust management should focus on designing and implementing architectures for collecting trust relevant information and to derive measures of trust and reputation. Once that is done, these measures can be integrated into classical economic models.

This paper therefore focuses on the architectures and mechanisms for reputation systems, and not how they are to be integrated into decision models.

The concept of trustworthiness is closely linked to that of reputation, but we would like to point out that there is an important difference. For the purpose of this study we will utilise and accept the definition of reputation given by the Concise Oxford Dictionary.

Reputation:
Reputation is what is generally said or believed about a person's or thing's character or standing.[8]

This definition corresponds well with the view of social network researchers [9,10] that reputation is a quality derived from the underlying social network which is globally visible to all members of the network. The difference between trust and reputation can be illustrated by the following perfectly normal and plausible statements:

1. "I trust you because of your good reputation."
2. "I trust you despite your bad reputation."

Statement (1) reflects that the relying party is aware of the trustee's reputation and bases his trust on that. Statement (2) reflects that the relying party has some private knowledge about the trustee, e.g. through direct experience or intimate relationship, and that these factors overrule any reputation that a person might have. This observation reflects that trust ultimately is a personal and subjective phenomenon that is based on various factors or evidence, and that some of those factors carry more weight than others. Personal experience typically carries more weight than recommendations from others, but in the absence of personal experience trust often has to be based on recommendations from others.

Reputation can be considered as a collective measure of trustworthiness based on the opinions of members in a community. An individual's subjective trust can be derived from a combination of received recommendations and personal experience.

Reputation can refer to a group or to an individual. A group's reputation can for example be modelled as the average of all its members' individual reputations, or as the average of how the group is perceived as a whole by external parties. Tadelis' [11] study shows that an individual belonging to a given group will inherit an *a priori* reputation based on that group's reputation. If the group is reputable all its individual members will *a priori* be perceived as reputable and vice versa.

3. Applicability Criteria for Trust and Reputation Systems

This section describes essential elements for the applicability of trust and reputation systems in a given context. In order to illustrate these elements the eBay (http://ebay.com) reputation system will be used as example.

3.1. Basic Trust Dimensions

Trust requires a relationship between two parties with a specific scope. For any potential application the three basic elements: *originator*, *target* and *scope* of trust [12] must be implicitly or explicitly expressed..

3.1.1. Trust Originator

The trust originator represents the *trusting party* or the *relying party* in a trust relationship. In other words the originator is the party that enters into a situation of dependence. For example, on eBay, the auction bidders represent trust originators.

When assuming that trust is only meaningful to conscious agents the relying party is normally a human. Through a slight abuse of language it is commonly said that an organisation can trust another party. In reality certain individuals in that organisation trusts the other party and their opinion is taken as the opinion of the whole organisation.

When assuming that systems such as e.g. software agents can act on behalf of their human masters, it can also be said that systems can trust. The trust perception that can be given to systems will only be as complete or correct as the trust model and its implementation within those systems.

3.1.2. Trust Target

The trust target represents the *trusted party* on which the relying party depends. The trusted party can be almost anything. Humans and human organisations can of course be trusted, but so can systems, items and even abstract notions whenever a trusting party depends on it. To use eBay as an example, the sellers represent the trust target.

In order to have more than a baseline trust in the target entity it must be possible to identify or recognise the entity. A potentially complicating factor is the under-specified concept of identity [13]. This problem can be formulated with the following two questions:

1. *Are the things we would like to regard as identifiable entities, really entities that have identifiers?*
2. *Does the relationship between identifiers and identifiable entities have properties we can make maximal use of?*

The answers to these questions will depend on the particular context. In general it must be possible to answer both questions in a clear and positive way in order to have a suitable trust target.

3.1.3. Trust Scope

The scope represents the operational aspects of what the originator expects from the target when entering into a situation of dependence. In contrast to defining the trust originator and target, it can be challenging to define the trust scope precisely. For example, on eBay, the trust scope is quite general and poorly defined. It can best be described as a composite trust scope consisting of correct product representation, good product quality, prompt delivery, and fair refund practice.

3.2. Trust Relationship Classes and Instance

Within each industry sector one or several groups of entities can play the role of trust originator and trust target, and there are multiple types of possible trust relationships between each pair of groups. In the health sector, a target group can for example consist of surgeons, and a set of trust scopes can for example relate to the various qualities that make a surgeon trustworthy. Honesty and skill can for example be two different trust scopes in that regard, and this could form a set of trust scopes. A given pair of originator and target groups connected by a given set of related trust scopes can be called a *trust relationship class*.

Within a given trust relationship class the originator and target groups can contain an arbitrary number of individual entities, and each set of trust scopes can consist of one or several specific trust scopes. We will use the term *trust relationship instance* to denote an actual trust relationship involving a specific originator, a specific target and a specific trust scope. The multidimensional aspects of trust relationships can thus be recognised both on a class level and on an instance level as illustrated in Fig.1 below.
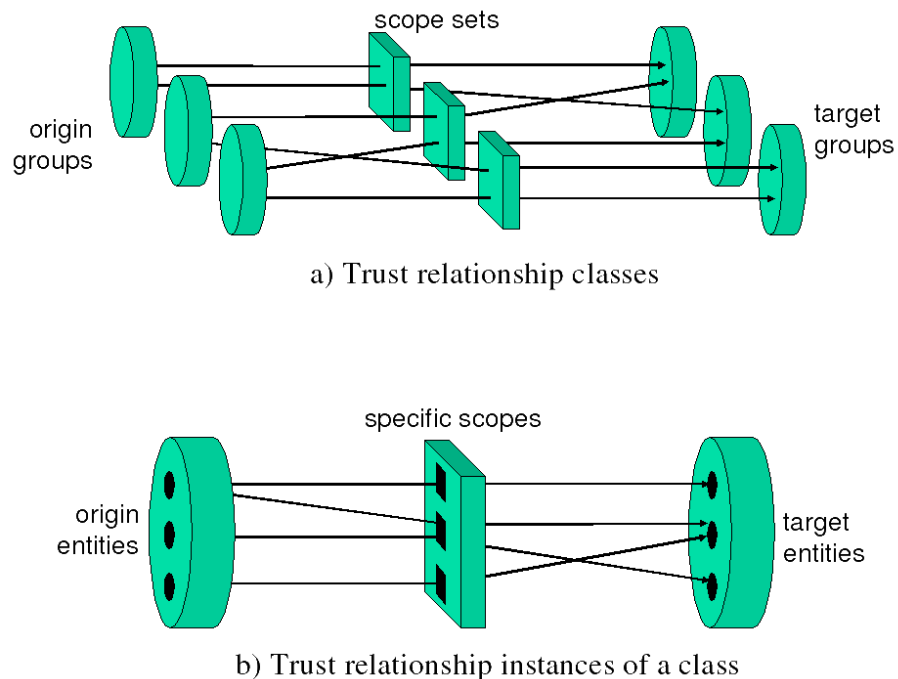

a) Trust relationship classes

b) Trust relationship instances of a class

Figure 1. Multi-dimensional trust relationships

The arrows in Fig.1.a represent relationship classes, and the arrows in Fig.1.b represent trust relationship instances. Given the large number of possible trust relationship classes as well as the large number of possible trust relationship instances within each class, the theoretic number of possible trust relationships is huge. In practice however only a relatively small subset of all possible trust relationship classes will be relevant.

3.3. Conditions for Trust

The situation of dependence must include risk and uncertainty. More specifically there must be possible harm to the relying party, and the trustworthiness of the trusted party must in general be subject to varying degrees of certainty.

3.3.1. Dependence

The relying party must be in a situation of dependence on the trusted party for a particular trust scope. For example on eBay, the buyer pays before receiving the goods, and depends on correct product representation, on good delivery, and finally on fair refund practice in case something goes wrong.

3.3.2. Potential Harm

In case the expectations of the trust scope are not met by the trusting party the relying party will suffer some form of harm. The more precisely the potential harm can be identified and quantified, the easier it is to use reputation systems as decision support tools. The potential harm to the buyer on eBay is to loose money.

3.3.3. Uncertainty

In order for reputation systems to be useful, the trustworthiness of the trusted party must not always be known by the relying party because this is precisely the type of information a reputation system is designed to provide. In other words it must be plausible that the trust target behaves or performs contrary to expectations. In the case of eBay, without the reputation system, the buyer would have no information about the past performance of the seller, and would therefore be very uncertain.

3.4. Input and Communication Requirements

It must be possible to obtain suitable information that can be fed into a reputation system. This requires that the information can be discovered and communicated.

3.4.1. Input Requirements

The information needed as input to reputation systems can be represented in various forms from verbal and qualitative to formal and quantitative. It must be possible to use this information to derive measures of trustworthiness or reputation which also can be represented in various forms. It must be possible to determine where this information can be found or how it can be extracted. The information must also be in a form that is suitable for processing. For example on eBay, ratings are received from the participants in the form of positive (+1), negative (-1) and neutral (0) ratings. The computation principle used by eBay is simple summation of all received ratings.

3.4.2. Communication Requirements

Given that it is possible to locate or derive information it must be possible to transfer this information to where measures of trust or reputation are derived. The derivation can be done by a centralised party and the result communicated to the relying party, or the derivation can be done directly by the relying party. In practice it must be possible to design communication protocols to transfer the input information for the reputation systems and the derivation output results. The reputation system on eBay is centralised because eBay receives all the ratings and derives the reputation score.

4. Approach

This section outlines a general approach for identifying potential applications of reputation systems within an industry sector. It mainly consists of identifying elements in the categories described in Sec.3 and assessing their suitability for being formalised into reputation systems. Because the applicability of a reputation system to a given category of trust relationship depends on all associated elements, a partially integrated approach should be taken. In practice this would take place by identifying relevant trust relationship classes on an intuitive basis, while at the same time try to determine whether the trust aspects of dependence, risk and uncertainty are present. The last part consists of determining the most suitable reputation system mechanisms.

For each industry sector we will identify groups that can play the role of trust originator and trust target, and try to determine typical trust scopes that can form part of trust relationships between these groups. Not all combinations of originator/target/scope will be meaningful, and simple intuitive judgement should be used to identity the most relevant combinations.

For a given trust relationship it must be possible to identify the existence of uncertainty and possible harm. The uncertainty must be in the eyes of the relying party and must relate to the performance of the trusted party. In case the target is a conscious agent this assumes an information asymmetry between the originator and target groups in that the relying party possesses less information about the honesty and/or reliability of the trusted party than the trusted party itself. The purpose of the reputation system will then be to collect as much information as possible and use this as the basis to derive trust in the target. In case a quantitative assessment of possible harm is required a process similar to risk analysis might be needed.

After having identified relevant trust relationship classes, the practical issues related to communication protocols and computational engines can be worked out. For this it is necessary to consider how information will be collected and how measures of trust can be derived. At this stage, usability aspects can be taken into considerations. The success of a reputation system depends on how the human users respond

to it. While usability principles can be used in the design phase, experience from practical implementations is the only way to verify that usability requirements are satisfied. The purpose of this paper is not to describe a reputation system for the health sector on a deep technical level, but rather to provide a general approach.

5. Reputation Systems for the Health Sector

5.1. Trust Issues in the Health Sector

The health sector tries to enforce very high standards of care and treatment of patients. This requires for example highly skilled health professionals, efficient administration procedures and high quality drugs.

Although it is in everybody's interest to have the best possible health system, there are situations where individuals and organisations could profit from compromising quality. Threats to the health system could for example emerge in the form of any medical staff trying to hide the fact that their skills do not meet the standards, or administrative staff who are unable or unwilling to improve efficiency due to lack of administrative skills, or in order to protect their jobs. One example from Australia is the relatively recent case of Dr. Jayant Patel who previously had been placed on probation for 3 years in 1983 for "gross negligence" in his practice at Rochester Hospital in New York State, whose scope of surgery had been restricted by the Oregon Board of Medical Examiners in 2000, and who had been threatened of having his licence revoked in New York State in 2001, which he avoided by obtaining permission to surrender his licence to practise. It is only by refraining from informing Australian health authorities about these incidents in the USA that he was able to get promoted to Director of Surgery at Bundaberg Hospital in Queensland, Australia in 2003. The subsequent questioning of Patel's performance at Bundaberg Hospital did not emerge from a clinical governance system but from concerns of individual doctors and nurses about his surgical performance and prowess. It was a letter from the nursing staff about this matter which, when tabled in Queensland Parliament, resulted in the establishment of a Commission of Inquiry headed by Anthony Morris QC. In the meantime, Patel left Australia unimpeded. [14]

While it is impossible to avoid all medical errors, there must be rigorous standards in place, and procedures for implementing changes and improvements when things go wrong. It is to be expected that treatment quality and cost/efficiency are in inverse proportion to each other. By accepting this fact it would be useful to be able to measure both in order to identify those health institutions that are able to find the most optimal balance. Reputation systems could be used to assist in achieving an optimal balance.

5.2. Existing Applications

Health professionals traditionally have very strong organisations protecting their interests. Two important aspects of this are to protect members from external criticism, and to enforce professional standards internally. Health organisations such as the Australian Medical Association (http://www.ama.com.au/) have procedures for sanctioning members who breach ethical standards or commit professional negligence. Although these procedures operate with little automation, they could be called reputation systems.

Hospital performance is usually measured as a function of the number of patients they treat. However, this is a quantitative measure that does not give any indication of patient satisfaction.

5.3. Health Reputation System Scenarios

5.3.1. Surgeon Performance

Trust originator:	Patients
Trust Target:	Surgeons
Trust Scope:	To perform successful surgeries

Table 1. Trust relationship class: Surgeon performance

A consequence of insufficient reporting of medical negligence is that the public often has very limited access to information related to breach of ethical standards and professional negligence. For example, there is significant deviation between Australia and other countries regarding the legal obligation of hospitals to report such cases [15]. In the United States there is a National Practitioner Databank which keeps a wealth of information on medical practitioners including: disciplinary actions taken against a practitioner; civil suits filed against the practitioner for malpractice; any downgrading of hospital privileges or dismissals from employment etc. The Databank is a commercially run National Government Organisation which charges for information. Any current or prospective employer or registration board may apply for a doctor's complete record. However, the public does not have access to this Databank. However, there are state-based initiatives in the United States to publish complaint and litigation data about doctors on the web. The Massachusetts Board of Registration in Medicine has been at the forefront with this issue. The argument in support of this has been that health consumers should be able to make an informed choice about which doctor they wish to attend. In Australia, cases of medical negligence are reported to state government bodies such as the Health Care Complaints Commission (HCCC) in NSW, but not to the public. A reputation system based on direct input from patients would represent an addition a substitute for inadequate reporting. In case public reporting is implemented, such as in Massachusetts in the USA, a reputation system would provide the public with additional information about health practitioners, and would for example provide a method for allowing high professional quality to be recognised.

A nationwide centralised architecture together with Bayesian computation would be a possible way of implementing this type of reputation system, but it could also be implemented in a more distributed fashion, e.g. per hospital.

Aggregated scores of surgeons from a given hospital can be used to measure the quality of the hospital from a surgery perspective. Measures could also be broken down into surgery categories. Those measures could also be compared with the operational costs in order to derive a measure of cost/quality effectiveness.

5.3.2. Drug Effectiveness

Trust originator:	Patients
Trust Target:	Drugs
Trust Scope:	To be effective, and not have undesirable side effects

Table 2. Trust relationship class: Drug Effectiveness

In most developed countries, including Australia, drugs are subject to strict safety trials and analyses before they are released into the market. Drug evaluation and approval is for example done by the FDA (Food and Drug Administration) in the USA, and by the TGA (Therapeutic Goods Administration) in Australia. Subsequent to approval the drug companies often run aggressive promotion campaigns using the same methods that are being used to promote any other consumer items. These campaigns are partly targeted directly at patients/consumers and partly at medical professionals who prescribe the drugs to the patients. There is a danger that prescription of certain drugs is biased by explicit or implicit benefits offered by the drug companies to those who prescribe the drugs. A comprehensive report by the Health Committee of the UK House of Commons (22 March 2005) describes how drugs are "intensely promoted to prescribers" including the promotion of unsafe drugs [16][4]. A report by the West Virginia Pharmaceutical Cost Management Council [17][5] found that drug companies spent $6000 to $7000 per doctor on marketing to physicians. The report also describes how Drug companies use drug reps in encouraging physicians to switch patients to their company's drugs, which are often more expensive than generic drugs or drugs offered by competitor firms, and that Drug companies use the practice of giving gifts and financial support to physicians and prescribers,

The process above is poorly suited to select drugs as a function of patients' actual satisfaction with taking particular drugs. Medical professionals try to observe the effect of drugs on an objective way, but they usually do not take notice and report the patients' subjective satisfaction to the drug companies. From the

---

[4] http://www.parliament.the-stationery-office.co.uk/pa/cm200405/cmselect/cmhealth/42/42.pdf

[5] http://www.state.wv.us/got/pharmacycouncil/Advertising_Marketing_PrescriptionDrugs.pdf

patients point of view both aspects are important, and assuming that two drugs have the same observable effects, the one with the highest patient satisfaction would be expected to be preferred. Additionally there are many aspects of drugs that are difficult to measure in an objective way such as how convenient they are to use, and how inconvenient their side effects are.

Drug efficiency could be measured by those who are directly exposed to them, namely the patients. Electronic medical records that allow patient input can provide an efficient way of collecting and aggregating information about the objectively observable effect of drugs, and of the patients' subjective satisfaction. In order to allow patients to provide subjective ratings about drugs, a nationwide centralised architecture together with Bayesian computation could be used for implementing this type of reputation system. It could also be implemented in conjunction with specific eHealth initiatives such as the Australian HealtConnect[6].

5.4. Implications of Reputation System in the Health Sector

The health sector consists of powerful groups such as professional associations, the drug manufacturers and the hospitals. The patients are relatively deprived of influence in this constellation. Reputation systems have the potential of changing this balance by increasing the patients' power. It is therefore to be expected that reputation systems will be met with resistance from some or all of these groups who have a vested interests in the status quo. In case of resistance from the health profession, the reputation systems described in the scenarios above can operate without cooperation from those groups, and can be driven by patient interest groups. In the other hand, it it may also be possible to get for example the drug manufacturers to cooperate if reputation systems can create communities where patients can obtain additional information about drugs from the drug companies similarly to how product vendors support the community around the product rating website Epinions.com[7].

However it seems less likely that surgeons or hospitals would cooperate with setting up and running reputation systems if they will be subject to patient ratings. The internal procedures they already use give them better control of how information is published. The problem of defamation and possible misuse of reputation systems must be reduced to an acceptable level by implementing robust control mechanisms around the reputation systems. The reputation systems must be trusted by all parties to produce reliable and representative information about drugs, health providers and practitioners.

6. Reputation Network Architectures

The network architecture determines how ratings and reputation scores are communicated between participants in a reputation system. In centralised reputation systems, a reputation centre collects information about participants in the market or community. The reputation centre typically derives a reputation score for every participant and makes all scores publicly available. Participants can then use each other's scores for example in decision making. The idea is that transactions with the most reputable participants are likely to produce favourable outcomes. In a market with strictly symmetric relationships, every member can rate every other member. In a market with asymmetric relationships, such as where there are well defined groups of service providers and service consumers, it is normal that members of one group rate members of the other group. P2P networks are symmetric markets, whereas the health industry typically has asymmetric markets. The medical doctors and their patients represent such a market, where the medical doctors and the patients form separate groups within the market. It is of course possible for individuals to be members of both groups, but they normally only play one of the roles in any specific interaction.

Fig.2 below shows a typical centralised reputation framework where C represents service consumers, and P denote service providers. In Fig.2.b, the service consumer considers using a service from the provider based on the experience of others in the past.

---

[6] http://www.health.gov.au/internet/hconnect/publishing.nsf/Content/home
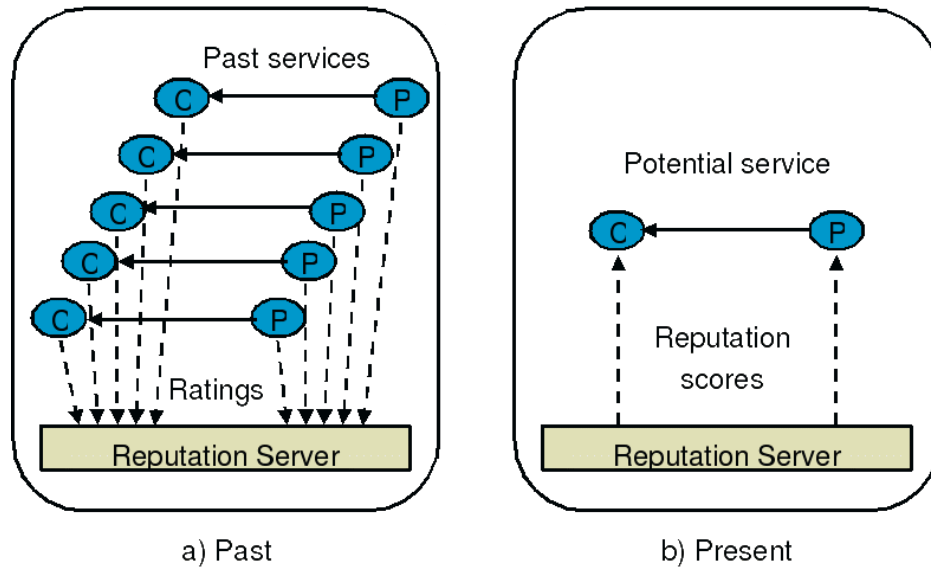[7] www.epinions.com

Figure 2. General framework for centralised reputation systems

After a transaction is completed, the consumers provide ratings about the providers' performance. The reputation server collects ratings from all the consumers and continuously updates each provider's reputation score as a function of the received ratings. Updated reputation scores are provided online for all the consumers (and providers) to see, and can be used by consumers to decide whether or not to use services from a particular provider.

The two fundamental aspects of centralised reputation systems are:

1. *Centralised communication protocols that allow participants to provide ratings about transaction partners to the central authority, as well as to obtain reputation scores of potential transaction partners from the central authority.*

2. *A reputation computation principle} used by the central authority to derive reputation scores for each participant based on received ratings and possibly on other information.*

One way of implementing an online reputation system for health services could be to create an online directory of the various services. This can be similar to the Sensis Yellow pages service directory. In addition, the directory service must be linked with the actual reputation service, so that reputation scores can be provided together with search results. This architecture is illustrated in Fig.3.
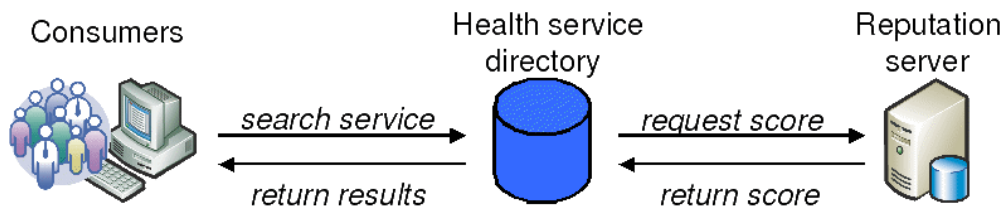

Figure 3. Reputation system architecture for fetching reputation scores

The service providers must be able to restrict ratings to only be given by consumers who have actually received health services. This can be achieved by the health service providers giving a one-time-ticket to a

consumer each time they have received a health service. By using cryptographic methods, the ticket can be made anonymous, so that consumers are able to rate services anonymously, without fear of any form of retaliation from health service providers that are rated negatively. This architecture is illustrated on Fig.4. This would prevent users who have not received tickets from providing ratings. Cryptographic verification of anonymous identities would prevent service providers from issuing tickets to fictitious users.
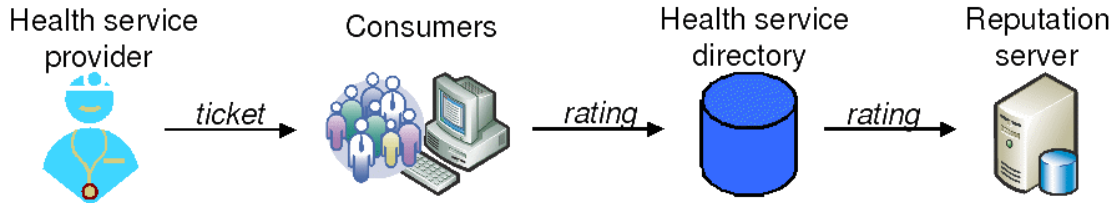

Figure 4. Architecture for providing ratings

In addition, the reputation system must be robust in the sense that it can not be misused e.g. by ballot-box-stuffing or defaming a particular health service provider. The robustness must be based on authentication and auditing of users and their behaviour. Without going into detail, here must be policies in place for how ratings can be provided, and mechanisms for sanctioning service providers and users who do not abide by these policies.

7. Discussion and Conclusion

Health service consumers find themselves in a wide range of situations of dependence where trust becomes important. Trying to formalise these situations for the purpose of building reputation systems can be challenging, and the approach described in this document consists of splitting the problem into smaller parts. Firstly, it must be possible to identify the groups that are involved in a trust relationship, and its trust scope, where the groups and the scope together form a trust relationship class. It is then useful to assess the aspects of trust such as dependence, potential harm and uncertainty relative to this trust relationship class. The idea is that a reputation system will only be useful if the uncertainty and the potential harm are significant. Finally it must be possible to extract or obtain trust related information about the trusted party for the purpose of computing measures of trust or reputation. This information can be provided by participants as subjective ratings, or objectively determined by observing or monitoring actual events.

This study shows that robust reputation systems can be applied to health service provision, and indicates a typical architecture for implementing such systems. Other architectures are also possible, and the choice of architecture must be made with the aim to provide a practical and efficient way of communicating this information to where the trust derivation will take place. The most suitable computational principle will depend on the type of information, on what the reputation system is supposed to measure (i.e. the trust scope), and on the usability requirements.

References

1. Doctor rating websites should be closed down: AMAQ. *ABC News Online*. 11 July, 2007 8:14am AEST. http://www.abc.net.au/news/stories/2007/07/11/1975394.htm. (accessed 30 August 2007).

2. A. Jøsang, C. Keser, and T. Dimitrakos. Can We Manage Trust?. *Proceedings of the Third International Conference on Trust Management (iTrust'05)*, 2005.

3. D.H. McKnight and N.L. Chervany. The Meanings of Trust. *University of Minnesota, Management Information Systems Research Center Working Paper Series 96-04*, (http://www.misrc.umn.edu/wpaper/wp96-04.htm), 1996.

4. D. Gambetta. Can We Trust Trust? In: D. Gambetta, editor. *Trust: Making and Breaking Cooperative Relations*. Basil Blackwell. Oxford. 1990, pages 213-238.

5. R. Falcone and C. Castelfranchi. Social Trust: A Cognitive Approach. In: C. Castelfranchi and Y.H. Tan, editors,

Trust and Deception in Virtual Societies. Kluwer 2001, pages 55-99.

6. A. Jøsang and S. Lo Presti. Analysing the Relationship Between Risk and Trust. In: T. Dimitrakos, editor. *Proceedings of the Second International Conference on Trust Management (iTrust)*. Oxford, March, 2004.

7. O.E. Williamson. Calculativeness, Trust and Economic Organization. *Journal of Law and Economics*, 36:453-486, April 1993.

8. *The Concise Oxford Dictionary, 8th edition*. Oxford University Press, 1990.

9. L.C. Freeman. Centrality in Social Networks. In: *Social Networks*, 1:215-239, 1979.

10. P.V. Marsden and N. Lin. *Social Structure and Network Analysis*. Beverly Hills: Sage Publications, 1982.

11. S. Tadelis. Firm Reputation with Hidden Information. *Economic Theory* 21(2):635-651, 2003.

12. A. Jøsang A, E. Gray and M. Kinateder. Simplification and Analysis of Transitive Trust Networks. In: *Web Intelligence and Agent Systems)*, 4(2):139-161, 2006.

13. H. Anderson *et al*. Survey of Privacy and Information Technology (Chapter on Trust Reputation and Privacy). *SAITS Project Report, OrbiTeam BSCW* (http://bscw.sics.se/pub/bscw.cgi/0/162845), 2003.

14. M.B. Van Der Weyden. The Bundaberg Hospital scandal: the need for reform in Queensland and beyond. *The Medical Journal of Australia.* 183 (6): 284-285, 2005.

15. Committee on the Health Care Complaints Commission. *Report on Mandatory Reporting of Medical Negligence*. Parliament of New South Wales, November 2000.

16. Health Committee of the UK House of Commons. *The Influence of the Pharmaceutical Industry*. Fourth Report of Session 2004–05, 22 March 2005.

17. West Virginia Pharmaceutical Cost Management Council. *Advertising and Marketing of Prescription Drugs*. Report, 8 September 2005.

Correspondence

Dr. Audun Jøsang
  Email : a.josang@qut.edu.au