

Open Research Online

The Open University's repository of research publications and other research outputs

Trust Strategies for the Semantic Web

Conference or Workshop Item

How to cite:

O'Hara, Kieron; Alani, Harith; Kalfoglou, Yannis and Shadbolt, Nigel (2004). Trust Strategies for the Semantic Web. In: Workshop on Trust, Security, and Reputation on the Semantic Web, 3rd International Semantic Web Conference (ISWC'04), 7-11 Nov 2004, Hiroshima, Japan.

For guidance on citations see [FAQs](#).

© 2004 The Authors

Version: Accepted Manuscript

Link(s) to article on publisher's website:

<http://iswc2004.semanticweb.org/>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Trust Strategies for the Semantic Web

Kieron O'Hara, Harith Alani, Yannis Kalfoglou, and Nigel Shadbolt

Intelligence, Agents, Multimedia Group, School of Electronics and Computer Science, University of Southampton, UK
{kmo, ha, y.kalfoglou, nrs}@ecs.soton.ac.uk

Abstract. Enabling trust on the Semantic Web to ensure more efficient agent interaction is an important research topic. Current research on trust seems to focus on developing computational models, semantic representations, inference techniques, etc. However, little attention has been given to the plausible trust strategies or tactics that an agent can follow when interacting with other agents on the Semantic Web. In this paper we identify five most common strategies of trust and discuss their envisaged costs and benefits. The aim is to provide some guidelines to help system developers appreciate the risks and gains involved with each trust strategy.

1 Introduction

Trust is at the heart of the Semantic Web (SW) vision. Trust is a method of dealing with uncertainty; when dealing with independent agents, institutions or providers of resources (including knowledge), one trusts them if one accepts their characterisation of what they will do. Trust can be a moral notion (X trusts Y to act in X's interests), or not (X trusts Y to perform some task T). Adopting the attitude of trust towards others means that one can plan and cooperate more efficiently, at the cost of greater risk of wasting resources when trust is misplaced.

The SW, conceived as a *collection* of agents, will therefore function more effectively when trust is licensed. As pointed out in [2], trust is essential for agents collaboration; each agent will have to make subjective trust judgements about other agents with respect to the services they claim to be able to supply.

There has been a good deal of research in this field [19]. But from the SW point of view, there are some interesting higher-level problems still to be addressed [26]. In the first place, many of the approaches are firmly in the field of multi-agent systems, where trust is also clearly a deep issue. However this has meant that some of the important problems specific to the SW are inevitably glossed over. Furthermore, many approaches have been technical 'fixes' which make a good deal of sense in circumscribed, specific contexts, such as commercial negotiations exploiting some relatively rigidly defined notion of 'agent'.

Secondly, because of the technical nature of many solutions, there are few signs of consensus emerging in the field. In itself, this is not a serious problem - the more solutions, even partial ones, the better, for heterogeneous and distributed systems. But we argue in this paper that some higher-level patterns are emerging, and outlines of general approaches becoming visible, which may be of help in understanding and comparing current research, and planning future developments. These higher-level patterns can be understood as *strategies* for placing trust. Or, in other words, general attitudes towards other agents in conditions of uncertainty. The best strategy will vary based on many variables, such as context, risk, cost, task undertaken, etc.

2 Situating Trust for the Semantic Web

There are a number of challenges to set an infrastructure that could promote trust, not least the lack of consensus in the field as to the best way of doing it [26]. How should trust be modelled? What information *is* relevant? What features should online trust have? Is it analogous to offline trust? How do we preserve the link between trust of other agents, and the incentives for those agents to be trustworthy?

Let us briefly set out some of the more pertinent challenges. To begin with, the SW, as a collection of agents, will be highly distributed. In such a structure, the idea of a centralised authority is hard to sustain, though such authorities are useful for fostering trust. More plausible is the idea of various competing authorities in restricted domains. Against that is the idea of an agent being responsible for gathering enough information for its own trust judgements; this would be ideal in a large distributed system, but may be prone to error, depending on how many information sources there were and how reliable they were [6][4].

Second, assuming that trust on the SW is not totally centralised, an agent will have to be able to discover relevant information from a variety of heterogeneous sources. [11] suggests that an agent must combine information from various sources effectively (eg from its own experience and that of other agents, agents certificates, and agents's roles), estimate the trustworthiness of those sources, and manage to cope with strategies of lying agents.

Third, agents must be able to exchange information effectively [14]. In particular, it must be possible to bootstrap trust in a context before there have been enough transactions between the protagonists to make firm judgements.

Fourth, how can we model trust without undermining incentives for trustworthiness (if an agent knows what signals I look for when judging trust, that agent now only has an incentive to give out those signals, not to behave in a trustworthy manner)? What knowledge should agents try to gather? And what properties of trust need be taken into account?

Fifth, as we have conceived the SW, trust is subjective and dependent on context [4]; an agent may be trusted to do one task with one set of resources, yet not trusted to do a different task with different resources. Will it be possible to produce models of trust that respect this context-sensitivity without increasing computational overheads too greatly?

In the next two sections, we will look at the costs and benefits of certain strategies of placing trust, and consider which are well-placed to address the challenges we have outlined in this section.

3 Trust strategies for the Semantic Web

A strategy for trust is an attitude towards the relative costs and benefits of secured interaction, unsecured interaction or no interaction at all. For instance, a safety critical system may regard avoiding catastrophic failure as an important part of its function, even if the risk is very low. Hence it will be prepared to accept high costs for refusing cooperation with other agents where failure could lead to a catastrophe. On the other hand, an agent might accept a broker's suggestion of an unreliable SW service if the cost of failure was low.

There are many strategies possible for dealing with trust on the SW. Examination of the SW and related fields leads us to sketch five potential strategies, based on a rough division of approaches. We identify five basic strategies (fig. 1).




Fig.1. Five basic strategies from placing trust.

Much variation is possible within these strategies, and we do not claim that these are the only strategies available, or appropriate for the SW. Furthermore, it is possible to imagine systems that adopt a mixture of these strategies opportunistically, depending on what information is available to them at any moment, what risk they are prepared to bear, and what resources they have available to them. In other words, a system might change between strategies dynamically as circumstances (costs, risks, or the environment, for example) change. As an example, see [5], which switches between the first two of the strategies we outline, depending on its previous successes and failures with strangers. Different strategies get used as the system's understanding of the environment alters.

Note also that our discussion ignores the question of who takes the actual trust decisions. It is neutral on the question of how the preferences of people or policies of corporations interact with the agents' trust strategies. Clearly it is unlikely that software agents would get complete autonomy in important transactions.

3.1 Optimistic Systems:

Optimistic systems accept others unless there is reason not to trust. If the benefits of cooperation are relatively large or the costs of betrayal are relatively small, risk is low, and the gains from trust massively outweigh the gains from distrust. Such systems are also likely to benefit from decentralisation; self-policing ensures the shared ethics of the user population determine the policies of trust management [13].

Other systems exploit optimism for the purposes of a *bootstrapping process*. For example, in referral systems [27], each user has a personal agent which stores a user model. A user sends a query to his agent, which then suggests potential contacts to whom to send the query (which is an iterative process). In such a system users can form models of others on the basis of the replies they give. However, they need replies - and they need to act on them - if they are to gain any information that will enable them to adjust their acquaintance models. The system will bootstrap more quickly with an optimistic strategy.

Optimism Optimism is a very simple strategy. Basically it is the idea that an agent will trust another agent even if its performance is uncertain, unless there are positive reasons for not trusting it. The basic idea is that trust is the default attitude. Hence an agent will accept the bona fides of other agents offering services unless they fail one or more tests.

3.2 Pessimistic Systems:

CS AKTive Space (CAS) [23] is a SW application, which provides browsable representations of the discipline of Computer Science in the UK. The information is harvested off the

Web and stored by various techniques to acquire knowledge. The approach chosen by [23] is to rely on sites whose brands are trusted, for example the Web sites of computer science departments in the UK, thus increasing the chance of getting sound data. [6] attempts to understand the SW as an ecology of agents, rather than a big database, and suggests the use of authorised delegations, that is the granting of a right to one agent by another, which must meet various constraints before they can be accepted by other agents. It is claimed that such systems, which take a generally pessimistic view of other agents and knowledge sources, have clear advantages in the distribution of trusting decisions ([12], [4], and section 3.3 below). [18] describe a pessimistic approach, though intriguingly hybrid. They suggest taking in information from past interactions about both the confidence of an agent in other agents, as well as their confidence in this agent. Given a quantity of such information, agents can be ranked in terms of their reliability. Such systems in effect take a high rank as evidence of reason for trust; hence in such a system many trustworthy agents may fail to be trusted.

Pessimism Pessimistic strategies restrict interactions with agents unless there is a reason to trust them. Note that the pessimism corresponds to trust via personal acquaintance in the offline world, which is the basic model of trust (local trust, [17]). Such a model of trust is not often capable of supporting and underlying very complex societies [7].

3.3 Centralised Trust Systems:

A system which helps users annotate information sources is described in [8]. It provides them with a formalism for expressing agreement/disagreement, and the argumentative stance of the source. This is then used to measure a context-sensitive evaluation of the source. [15] proposes a centralised agent to measure the reputation of Web services by monitoring and collecting client feedback, and making this information available to other agents.

Relying on centralised institutions to measure trust takes the burden off the interactive agents when deciding which agents to trust. However, such systems raise the question of how trustworthy are the sources of their trust information in the first place, and why such trust warehouses should be trusted at all [4]. One observation made in [20] with respect to *eBay* is that users feedback is almost always positive. The authors note that most people do not like giving negative feedback, unless revenge is a motivation. Clearly the centralised trust system, be it *eBay* or otherwise, is only as good as the information it is receiving and brokering; if, as with *eBay*, there is a bias in the information coming into the system and being displayed, then the strategy of other agents of shifting trust to the centralised system may well be non-optimal. Similarly, [6] argue against centralised units for measuring trust because of their scalability limitations and the implicit trust measurement mechanisms they adopt.

Centralisation Centralising trust involves laying off the costs of interacting with and investigating agents to a central institution or authority. The institution may certify a particular agent as willing and able to act in some specified way. If the agent bears a certificate, then it could be trusted. However, this does not obviate the need for trust, but the trust requirements are reduced. The agent now only needs to trust the institution - can it investigate potential agents thoroughly, and does it have the powers to sanction betrayals effectively? Some institution, like [8]'s TRELLIS system, merely holds the relevant collected information for users to process themselves.

3.4 Trust Investigation Systems:

In [24], a Bayesian system is describe which models trust in a P2P file distribution network. On such a network, peers make recommendations to each other about where suitable files might be found. The agents 'gossip' with each other, by exchanging and comparing their Bayesian networks. After this comparison, the agents update their trust ratings of each other, depending on whether they share similar preferences, on the assumption that an agent with similar preferences is more likely to give suitable recommendations than others.

In other words, the agents perform an investigation of the others in order to determine how likely it is that their recommendations will be useful. The agents do not simply receive the recommendations passively; they compare Bayesian networks, and undertake extra computational tasks, in order to determine their suitability.

Another example of this sort of approach is provided by systems that negotiate automatically to extract trust credentials from other parties. For example, the TrustBuilder system [25] works by iterative disclosure of credentials by the negotiating parties, in response to a series of requests for credentials, until either the set of credentials is exhausted, and the negotiation concludes unsuccessfully, or until sufficient trust between the parties has been established. The credentials are usually supplied by centralised authorities, and so this version of investigation has strong links with the previous strategy. But the point of a negotiation is that the agents themselves do (potentially a lot of) computation on the credentials, the context and the current state of negotiation to work out (a) which credentials to ask for, (b) how to evaluate responses, and (c) how much information to release at each stage in the negotiation.

Investigation This suggests a fourth strategy. Trust is a response to uncertainty. But trust imposes risks. Hence, to avoid some risk, one strategy is to reduce uncertainty by investigating or evaluating other agents to determine some salient details of operation. It is not passive; it actively tries to discover aspects of the environment that are relevant to reduce uncertainty.

3.5 Transitive Trust Systems:

In [10], authors argue that a likely future scenario will be a networked infrastructure of entities, linked by various ad hoc networking technologies. They use the small world theory [16], which hypothesises that any pair of objects in a random network will be connected by a relatively short chain of random acquaintances. This entails that if such mutual chains of acquaintances are used to determine initial trust between a pair of entities, then the method will scale up well because these chains are likely to be small.

Social network analysis techniques are used in [9] to measure trust over a Friend of a Friend (FOAF)¹ network, extended with trust relations. A similar approach of exploring webs of trust is described in [21], where users provide trust values for a number of other users, which are used to measure trust. These algorithms, as in [8], rely on manually-set trust opinions, which once more shifts the problem of trust to the source of the opinions.

As [3] are correct to argue, trust is not strictly transitive. If A trusts B, and B trusts (and maybe recommends) C, nothing follows about whether A trusts C. However, there are some circumstances where transitivity happens, at least locally. Similarly, transitive effects can be seen where trust is given a recursive definition [4].

¹ <http://www.foaf-project.org/>

One weakness of much of the work above is the lack of sensitivity to *context*. Context is a basic feature of trust [4] and must therefore be considered when dealing with it. For example you may trust a colleague to write a good project proposal, but you might not trust him to fix your car. Associating trust measurement with a specific context will inevitably increase complexity, but nevertheless is crucial for deriving meaningful trust values. Some notion of context could be obtained from the type of relations between networked agents, which plays a crucial role when measuring agent's reputation [22].

Exploring Transitivity These systems can be described as using the strategy of exploiting *transitivity*. The idea of this strategy is that an agent sends a message out about whether a potential agent is trustworthy. The network of acquaintances of that agent will then either send back an opinion based on experience, or pass the message onto its acquaintances, many of which will be unknown to the first agent. The aim is to increase the scope of an agent's knowledge by exploring the network feature of agents communities to bring in information from other, unknown, agents.

As noted above, the issue of context is important, and the trade-off between investigating context and relying on others' reports and recommendations shows the close relationship between the strategies of Investigation and Exploring Transitivity. In many actual systems and situations, there is likely to be a tight coupling between the two strategies, as investigation ameliorates some of the difficulties of exploiting transitivity.

4 Costs and Benefits

The main point of a trust strategy is to provide a way to operate under uncertainty, not taking too many risks, not missing too many opportunities, not deliberating too long before making commitments. Therefore, before setting up a plan of interaction between SW agents it is necessary to understand the risks and benefits associated with the selected trust strategy.

There are many different kinds of costs and benefits an agent might incur when communicating or dealing with other agents and services. Here we discuss four types of costs; *operational*, *opportunity*, *deficiency*, and *service charges*. We describe our estimations for each of these costs with respect to each of the five trust strategies discussed earlier, and summarise them in figure 2.

4.1 Operational cost

Operational costs are the expenses of operating a trust strategy. In other words, this is the cost of setting up and operating the whole trust plan. Therefore, the more complex the strategy is the higher the cost is expected to be.

The cost of operating an optimistic strategy is relatively small. Obviously it gets smaller the more optimistic the agent becomes, with a limiting case of zero when the agent takes all agents' *bona fides* as genuine. Pessimistic agents tend to run a set of tests to limit the number of agents they interact with. However, this does not necessarily lead to high operational costs. The one property that pessimism shares with optimism is that the strategy can be based on a very small and non-complex set of tests, which may involve merely the checking of identities. As the judgments about cooperation become more fine-grained, as for example with the investigation strategy, then operational costs will increase accordingly. One effect of trust is to cut out transaction costs, and naturally the effect of deciding to investigate rather than to trust is to reinstate them.

Agents in a centralised systems can look forward to very low operational costs because the complex parts of the strategy (e.g. investigations) will be handled by the institution itself (which can garner economies of scale). If the agent is part of a network where transitivity is adopted as a trust strategy, then it will incur some operational costs to implement and tune the techniques required to perform the network analysis. It may also incur costs as a by-product of being in a network, by having to pass on its experiences to other agents.

4.2 Opportunity cost

This is the cost of missing some possibility of generating benefit via interaction. The optimistic approach minimises opportunity costs. The more optimistic the agent, the less likely it is to miss some opportunity for interaction. On the other hand, there is the potential for a large opportunity cost with pessimistic agents by missing out on the potential to add value because of a reluctance to trust agents to perform necessary subtasks.

Opportunity costs may be higher with the centralised strategy than under the optimistic strategy as the agents' interaction will be limited to those certified by the institution. However, the cost should still be fairly low if many agents are certified. Opportunity costs are also likely to be high with the investigation strategy, unless the agent has the capacity and resources to investigate enough potential agents to keep opportunity costs down. With the transitivity strategy, the larger the network relative to the space of possible interactions, and the better its advice, the lower the opportunity costs.

4.3 Deficiency cost

Defecting agents fail to act as they claimed. The deficiency cost is the cost of betrayal by an agent. This cost is strongly associated with the amount of risk undertaken by an agent.

The risk of betrayal is high with optimistic agents, and so, the agent should plan for likely cost of an agent defecting. Equally for pessimistic agents the deficiency cost should be low. The main advantage of investigation is the lowering of risk by the simple expedient of lowering uncertainty. By performing an investigation or evaluation, the agent gains knowledge about whether and how the other agent is going to work.

The big advantage of a centralised system is that a small outlay of trust on the part of the agent still allows interaction with a large number of agents. However, this increases systemic risk. In a distributed system, a betrayal of trust will typically cause the agent to withdraw trust from the offending agent. But if an agent comes to another with a certificate issued by an institution, and then betrays that agent, the betrayed agent may well withdraw trust from that institution's procedures, and therefore refuse to cooperate with any of the institution's client base. The scalability problem of centralised systems adds to the systemic risk [4]. As the number of certified agents increases, then the degree of trust in the central authority may well begin to fall. Hence [4] recommend a strategy of pessimism, combined with a distributed system of trust management.

As for the transitivity strategy, if the network is large, then deficiency risk may also be large. But, unlike with an institution, it may be possible to weed out the source of bad advice from a network, and the network may be able to reconfigure itself and recover. Hence the advantage of exploring transitivity is that its costs may undercut those of the other strategies, and will also be spread, thereby hedging the risk.

4.4 Service charges

Agents may have to pay for purchasing services from other agents. Even on the most un-commercial views of the SW, there is little dispute that there will be a major effort required to populate it. Many people will of course be more than willing to post information on the SW, but the effort required to populate the space with enough RDF, and to ensure that the information posted will be that required or desired by other users will be non-trivial. Incentives of some kind will be required if the SW is to spread out of the academic ghetto.

Incentives and other charges may be high with the optimistic strategy for two reasons. First, the optimistic agent may well purchase more services from agents than those pursuing other strategies. Second, in the absence of strict monitoring of agent performance, higher payments may be required to ensure the alignment of interests.

Payment for services by pessimistic agents may be lower, as in some circumstances the agent can act on the understanding that its interests are aligned with those of the other agent. An additional advantage is that the smaller number of agents with whom the main agent interacts will reduce the complexity of any decision-making.

Adopting the investigation strategy may result in high operation costs, as its investigations themselves may well require consultation with many services. Similarly with the transitivity strategy - exploring the network may involve accessing fee-based services.

Operation costs of agents who rely on centralised units for handling trust is low. However, there might be a charge to be paid for this service. But because such institutions will probably serve many agents, they will benefit from economies of scale, and their charges should be much less than the accumulated costs of dealing with several distributed services.

Strategy	Optimism	Pessimism	Centralisation	Investigation	Transitivity
Operational cost	Low - does not require much policing.	Rises with complexity of filtering tests.	Low - the cost is embedded in the centralised service.	High - complex tests increase cost	Low if efficient methods are used.
Opportunity cost	Low - unlikely to miss many interactions.	High - does not check all possibilities.	High - decrease if many agents are certified.	High - only interact with investigated agents.	The larger the network, the lower the cost.
Risk	High - more risk when less cautious.	Low - won't interact with uncertain agents	Low - only legitimate agents are certified.	Risk is low by lowering uncertainty.	Rises with length of chains of referrals.
Deficiency cost	High - no check for malicious agents.	Low - malicious agents will be caught up front.	High - betrayal of one certified agent may affect the whole service.	Low - malicious agents banned if discovered	High - Behaviour of one agent might spread to network
Service payments	High - if max price/no. of interactions not set	Limited interactions results in limited charges.	Low - pay the centralised service only	High - may need to consult several services	May rise if when accessing fee-based services.

Fig. 2. Costs estimates for five trust strategies.

5 Meeting the Challenges of the SW

How do these strategies address the challenges for the SW set out in section 2? The cheap and simple strategies; optimism and pessimism, meet most of the challenges half way, by avoiding most of the complexities of the situation (eg. binary modelling of trust makes its computation much simpler). Bootstrapping is trivial with optimism, but very challenging with pessimism. Both approaches however, fail to take account of the context sensitivity of trust judgements.

Investigation as a strategy can be very useful for bootstrapping trust, as some relevant properties of an agent may be known before it has actually begun to operate in a domain. Furthermore, investigation could meet the context challenge, assuming the methods of investigation were sensitive to the changes in context. However, the distributivity of the SW will make thorough investigation increasingly costly.

Centralisation is similar in some respects to investigation. If centralised authorities distribute certificates of trustworthiness in some domains that increase in complexity, it may be hard for such authorities to scale up. That does not mean that this strategy should not be used; only that such systems can only expect to flourish with a certain amount of flexibility. Users of competing authorities will need to decide between them, which may impose transaction costs.

The most interesting systems are those that exploit transitivity, in that they allow relatively flexible responses. For instance, a system using the transitivity strategy could cope with the problems of scale caused by increasing distributivity by pruning its searches; as trust is affected by the length of a chain of recommendations, falling as the chain gets longer, this pruning is unlikely to result in the loss of too much investigation. In this way, the analysis of a social network is analogous to the discovery of communities of practice [1], where shorter paths indicate stronger links.

Where most transitivity systems break down most seriously is on context dependence. Systems such as those of [9][21] tend to rely on fixed trust values explicitly given by users for others to use. This is of course valuable information, but rather than respecting the subjectivity of trust, it creates an objective sum of others' subjective valuations - not the same thing. One could imagine network analysis techniques, on the lines of [1], being exploited here. If the agent could specify which relationships in a representation of a network are the most important for it, in its specific context, then trust values could be determined dynamically depending on that specification.

6 Conclusions and Future Work

The SW, will be large, heterogeneous, dynamic and uncertain. Trust will inevitably be an issue. But for trust to flourish in such an environment, agents must be flexible, and adaptive to new contexts. To do this, it is important to reason at a strategic level, to think in terms of classes of acceptable outcomes. To this end, an understanding of how the costs and benefits of various strategies vary across contexts is essential. As we have seen, different strategies can be advantageous in different circumstances. In general, given the SW challenges, the strategy of exploring transitivity by trying to create chains of trusting judgements, seems quite promising, though in certain circumstances other strategies will reveal their strengths.

This paper has produced a preliminary classification of strategies based on an a brief survey of approaches. This classification certainly is not the final word. But when a robust classification is available, it will remain to produce quantitative evidence about which strategies are valuable where - and when this evidence is available, it might be possible for actual trust mechanisms to be provided by SW service brokers. Such evidence might well be produced on experimental testbeds such as that described in [11].

Acknowledgments. This work is supported under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either express or implied, of the EPSRC or any other member of the AKT IRC. Thanks also to the anonymous referees for various helpful comments.

References

1. H. Alani, S. Dasmahapatra, K. O'Hara, and N. Shadbolt. Identifying Communities of Practice through Ontology Network Analysis. *IEEE Intelligent Systems*, 18(2): 18-25, 2003
2. T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, May 2001.
3. T. Dimitrakos, J. Bicarregui. Towards modelling e-trust. *Proc. 3rd Panhellenic Logic Symposium*, Anogia, Greece, 2001.
4. L. Ding, L. Zhou, T. Finin. Trust based knowledge outsourcing for semantic web agents. In *Proc. IEEE/WIC Int. Conf. on Web Intelligence*, Beijing, China, 2003.
5. M. Feldman, K. Lai, I. Stoica, J. Chuang. Robust incentive techniques for peer-to-peer networks. In *Proc. 5th ACM Conference on Electronic Commerce*, New York, 2004.
6. T. Finin, A. Joshi. Agents, trust and information access on the semantic web. *SIGMOD Record*, 31, 2002.
7. F. Fukuyama. *Trust: The Social Virtues and the Creation of Prosperity*. NY: Free Press, 1995.
8. Y. Gil, V. Ratnakar. Trusting Information Sources One Citizen at a Time. *Proc. 1st Int. Semantic Web Conf. (ISWC)*, Sardinia, Italy, 2002.
9. J. Golbeck, B. Parsia, J. Hendler. Trust Networks on the Semantic Web. *Proc. of Cooperative Intelligent Agents*, Helsinki, Finland, 2003.
10. E. Gray, J. Seigneur, Y. Chen, and C. Jensen. Trust propagation in small world. In *Proc. 1st Int. Conf. on Trust Management (iTrust'03)*, 2003.
11. D. Huynh, N.R. Jennings, N.R. Shadbolt. Developing an Integrated Trust and Reputation Model for Open Multi-Agent Systems. In *Proc. 3rd Int. Joint Conf. on Autonomous Agents and Multi-agent Systems (AAMAS)*, Columbia University in New York City, 2004.
12. L. Kagal, S. Cost, T. Finin, Y. Peng. A framework for distributed trust management. In *Proc. 2nd Workshop on Norms and Institutions in Multi Agent Systems*, Montreal, Canada, 2001.
13. S.D. Kamvar, M.T. Schlosser, H. Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. In *Proc. 12th Int. World Wide Web Conf.*, 2003.
14. B. Matthews and T. Dimitrakos. Deploying trust policies on the semantic web. In *Proc. 2nd Int. Conf. on Trust Management (iTrust'04)*, Oxford, UK, Mar. 2004.
15. E.M. Maximilien, M.P. Singh. An ontology for Web service ratings and reputations. *Workshop on Ontologies in Agent Systems, 2nd Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Melbourne, Australia, 2003.
16. S. Milgram. The small world problem. *Psychology Today*, 61, 1967.
17. K. O'Hara. *Trust: From Socrates to Spin*. Icon Books, Cambridge, 2004.
18. S.D. Ramchurn, N.R. Jennings, C. Sierra, L. Godo. A computational trust model for multi-agent interactions based on confidence and reputation. In *Proc. 2nd Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, Melbourne, Australia, 2003.
19. S.D. Ramchurn, D. Huynh, N.R. Jennings. Trust in Multi-Agent Systems. *The Knowledge Engineering Review*, in press, 2004.
20. P. Resnick and R. Zeckhauser. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In M.R. Baye (editor), *Advances in Applied Microelectronics*, vol. 11, Elsevier Science, Amsterdam, 2002.
21. M. Richardson, R. Agrawal, P. Domingos. Trust Management for the Semantic Web. *Proc. 2nd Int. Semantic Web Conf.*, Sanibel Island, FL, 2003.
22. J. Sabater, C. Sierra. Reputation and social network analysis in multi-agent systems. *First Int. Conf. on Autonomous Agents and Multiagent system*, Bologna, pp:475-482, 2002.
23. N.R. Shadbolt, m. schraefel, N. Gibbins, S. Harris. CS AKTive Space: or how we stopped worrying and learned to love the semantic web. *Proc. 2nd Int. Semantic Web Conf.*, FL, 2003.
24. Y. Wang, J. Vassileva. Bayesian network-based trust model. In *Proc. of the 6th Int. Workshop on Trust, Privacy, Deception and Fraud in Agent Systems. 2nd Int. Joint Conf. on Autonomous Agents and Multiagent Systems. (AAMAS)*, Melbourne, Australia, 2003.
25. M. Winslett, T. Yu, K.E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, L. Yu. Negotiating trust on the web. In *IEEE Internet Computing*, 6(6), 30-37, Nov/Dec 2002
26. Y. Kalfoglou, H. Alani, M. Schorlemmer, C. Walton. On the Emergent Semantic Web and Overlooked Issues. In *Proc. of the Third Int. Semantic Web Con. (ISWC)*, Hiroshima, Japan, 2004.
27. B. Yu, M.P. Singh. Searching social networks. In *Proc. 2nd Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, Melbourne, Australia, 2003.