Recommendation Privacy Protection in Trust-based Knowledge Sharing Network

Weisen Guo[1,2] and Steven Kraines[2]

[1] Inst. of Systems Eng. of Dalian University of Technology, Dalian 116024, China
guows@dlut.edu.cn
[2] Division of Project Coordination of the University of Tokyo, Tokyo 277-8568, Japan
{gws,sk}@cb.k.u-tokyo.ac.jp

Abstract. Trust can be applied to knowledge sharing on a distributed network of knowledge source agents. Each agent represents a person who trusts some other agents. Based on these trust-relationships, an agent can infer the trustworthiness of an unknown agent by asking trusted agents for recommendations. However, the person represented by an agent may not be willing to share his or her individual opinion about the trustworthiness of a particular agent to agents that do not protect information privacy. A solution for this issue is proposed using three kinds of privacy policies: generosity, caution, and non-cooperation. An agent that adopts the caution policy towards another agent will hide the details of the trust recommendation path. An analysis shows the effect of the privacy policies on the calculated reliabilities of the recommended trust values[3].

1 Introduction

Creation of semantic descriptions based on ontologies could be an effective way to represent knowledge that is better suited for computer-based processing and matching. We have described a four level framework for an agent-based scientific knowledge sharing network of independently distributed knowledge repositories on the Internet [1][2]. The second level prescribes tools for allowing experts to create computer-interpretable semantic representations of their knowledge content. We have developed EKOSS [3][4], the expert knowledge ontology-based semantic search system to implement this level. An agent-based network of distributed knowledge repositories would then be populated by multiple distributed and independent EKOSS systems acting as agents.

In the knowledge network, each agent represents a person sharing or seeking knowledge. When AgentA receives a recommendation for AgentC from AgentB, the recommendation will be the actual opinion of the person represented by AgentB, an inferred trust value calculated from other agent recommendations,

[3] Steven Kraines is the corresponding author. This work is supported by the Japan Science and Technology Agency through the Shippai Chishiki Project, the National Natural Science Foundation of China (70431001) and the Liaoning Province Science and Technology Agency (20061063).

or some combination of the two. In particular, if AgentA knows that AgentB is the final link of a trust chain to AgentC, then AgentA knows that the recommendation is the opinion of the person represented by AgentB because AgentB is not using any other information to infer the recommended trust value. However, the person represented by AgentB may not want his or her opinions to be public.

In our previous work [5], in order to calculate inferred trust values of unknown agents with a high degree of accuracy, all of the trust information along each trust path is returned to the source agent. However, in this case, the source agent will be able to determine how much the target agent is actually trusted by the person represented by the agent giving the trust value of the final link of a trust chain as explained above. This paper introduces a set of privacy policies into the trust model and evaluates the effectiveness of those policies to enable the agents to protect the privacy of the persons that they represent on the trust-based knowledge network.

The rest of the paper is organized as follows. Section 2 discusses a number of trust models in the literature that are related to our study. The privacy policy used by our approach is described in Section 3. In Section 4, several scenarios demonstrating the performance of the trust-recommendation network generated by the system are analyzed. The analysis shows how the effectiveness of the privacy protecting trust inferring algorithm depends on the situation of the trust network. Finally, in Section 5 we conclude the paper with a description of future work.

2 Related Work

Privacy has been a hot topic since as early as the 19th century, when an influential paper "The Right to Privacy" was published. Recently, the primary focus of privacy has shifted from media privacy, territorial privacy, communication privacy, bodily privacy, to information privacy as technologies for information sharing have progressed. The first four aspects of privacy have been well established in most legal frameworks around the world; however, information privacy continues to create many problems today [6].

Goecks and Mynatt [7] noted that privacy is a critical social issue confronting Ubiquitous Computing that requires urgent attention. They proposed that the concepts of trust and reputation are critical to understanding privacy and building systems that enable users to effectively manage privacy. They created a trust network that calculated the reputations of entities in the network. Based on these reputations, users can manage how, when, and where they share their personal information. Their approach offered a new way to protect the privacy of user's personal information and thereby addressed the problem of privacy protection in Ubiquitous Computing environments. We consider that a user's trust information for another user is his or her personal information. So, the privacy of user's trust information for others should be protected as well.

Our previous work introduced trust and recommendation concepts in a web-based knowledge sharing system. We presented the RTI algorithm to infer the trust value for the target agent from the recommendations of other intermediary agents in the trust chain to the target agent. The RTI algorithm uses negative as well as positive recommendations to accurately infer trust values of intermediary agents [5]. Most trust models do not try to evaluate the inferred trust value for the intermediary nodes in the trust chains [8]. However, if an agent, AgentA, is encountered in more than one chain from the source agent to the target agent, then ideally the source agent should give it the same trust value in each chain. The RTI algorithm infers the trust value for each intermediary agent in the trust chains to increase the accuracy of the inferred trust value for the target agent [5]. However, the RTI algorithm assumes that each agent will give all of the trust information that it has to a requesting agent. In this paper, we study the methods for protecting the privacy of trust information of individual agents.

3 Privacy Policy

In order to handle the issues that we presented in Section 1, we introduce a set of privacy policies in the RTI algorithm.

Like humans, when an agent receives a request for information on the trustworthiness of a target agent from a requesting agent, it should have the ability to decide what information it will return to the requesting agent. If the agent does not know or trust the requesting agent well, then it may not give any information to it. Even if the agent does know the requesting agent, if it cannot confirm that the requesting agent will protect its privacy, it may just return the information that it does not mind becoming public, such as an inferred trust recommendation for the target agent, and hide the detailed information of the recommendation chain. Only if the agent can confirm that the requesting agent will protect its privacy, will it return the trust recommendation together with all of the detailed information of the trust chain.

We provide three kinds of privacy policies to handle these three kinds of situations: generosity policy, caution policy, and non-cooperation policy.

When an agent AgentA receives a request for information on the trustworthiness of AgentB from AgentC and AgentA trusts that AgentC will protect its privacy by not giving the information to any unreliable agents, then AgentA will adopt the generosity policy and send to AgentC all of the trust recommendation information that it has that might be related to the trust chains to AgentB. If AgentA cannot confirm that AgentC will protect its privacy, but AgentA still wants to give some information to AgentC, then AgentA will adopt the caution policy and send only inferred trust values for AgentB, hiding the detailed information on the trust chain that led to the inference. If AgentA does not know or trust AgentC at all, then AgentA can adopt the "non-cooperation policy" and not give any recommendation information to AgentC.

Using the privacy policy, a person's privacy can be protected in the following way. As in our previous trust model, when the source agent wants to know the

trustworthiness of an unknown target agent, it sends out a request for trust information to the agents that it trusts, specifying a maximum chain length n (step 6 in our previous trust model [5]). If the receiving agent does not have a direct trust value for the unknown agent, it will send out another request for trust information to the agents that it trusts with a maximum chain length n-1. When an agent receives a request for trust information with a chain length of 1, it means that the requesting agent is asking only for direct trust values for the unknown agent, so any value that the receiving agent sends back will be understood to be the actual opinion of the person represented by the receiving agent. As we noted earlier, a person's actual opinions about the trustworthiness of other people is a form of private information. Therefore, an agent asked for trust information for an unknown agent with maximum chain length 1 will only return a trust value if it adopts the generosity policy towards the requesting agent (unless the agent being asked for information is a dishonest agent that is trying to trick the requesting agent with false information).

An agent that is requested for trust recommendations with a maximum length greater than 1 can return a trust value it has for the unknown agent even if it does not adopt the generosity policy. This is because the requesting agent cannot determine whether the recommendation is from the person represented by the agent or an inferred trust value calculated from recommendations of other agents, and so the actual opinions of the person represented by the agent are protected. Furthermore, in order to make full use of the RTI algorithm, an agent A can return the information that it has about an agent B between it and the target agent. However, if agent B is just on link away from the target agent, then the agent receiving the information about agent B from agent A will know the opinions of person represented by agent B. Therefore, agent A should only give this additional information to agents that it trusts highly, i.e. that it can guarantee to agent B to be trustworthy.

In the implementation that we are constructing, each person represented by an agent on the trust network has an interface to set the privacy policy adopted by her agent towards each agent that is known. The agent would initially adopt a default privacy policy, such as the caution policy. Later, the person represented by the agent could change the policy based on her assessment of the trustworthiness of the person represented by the target agent. Because each agent uses different privacy policies for engaging with both known and unknown agents, our modified trust system implements a form of basic privacy protection similar to real human interactions that should provide significantly more accurate trust inference than conventional systems based on statistical analysis of recommendations irrespective of source.

4 The Analysis of RTI algorithm with Privacy Protection

First, we revisit the scenario that we described in the previous paper, reproduced in Fig. 1 [5]. The scenario has a social network composed of ten users each characterized as having high reliability (H), moderate reliability (M), low

reliability (L), or as being dishonest (D). Sam is a dishonest service provider. For purposes of simulation, we assume without loss of generality that agents of users with high, moderate, and low reliabilities will give correct recommendations 90%, 80%, and 70% of the time, respectively, and the agent of a dishonest user will give opposite recommendations 90% of the time.

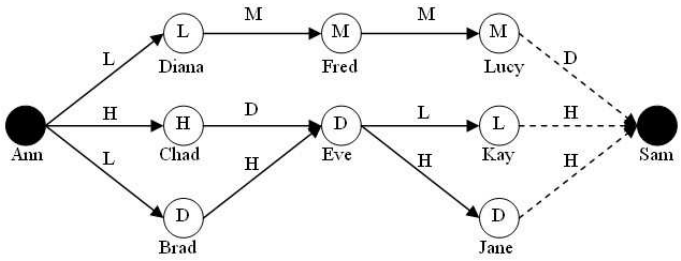

Fig. 1. The first scenario based on the scenario from [5]

The reliability of the service or recommendation trust value of an unknown agent can be calculated based on the recommendation trust values of the agents giving recommendations for the trust value of the unknown agent. We do this by giving each recommendation trust relationship a value between 0 and 1. Specifically, we give a low recommendation trust relationship a value of 0.7, a moderate trust relationship a value of 0.8, and a high trust relationship a value of 0.9. By quantifying the recommendation trust relationships in this way, we can combine recommendation trust values both in series (from the chain rule for Bayesian Networks) and in parallel (from the noisy-OR model for Bayesian Networks).

In this paper we add the privacy policies described in section 3 to the RTI algorithm, and we discuss the effect on the accuracy of the inferred trust values. If everyone adopts the generosity policy, Ann will receive all of the trust relationship information, including the identities and trust values of the target agents of each relationship. Therefore, Ann can calculate the inferred trust value of each inner agent in all trust chains and then use the RTI algorithm to obtain the most accurate inferred trust value for Sam.

If Brad and Chad both adopt the generosity policy towards Ann, then we have seen that the result will be the same as for the RTI algorithm with no privacy policy. If Brad adopts the caution policy towards Ann, and only Chad adopts the generosity policy, Ann will receive the information that Eve is dishonest from Chad and the recommendation that Sam is trustworthy from Brad. However, because Brad does not tell Ann that his recommendation came from Eve, Ann will use his recommendation, which she calculates as being more reliable than Diana's, and she will believe that Sam is highly trustworthy, which is incorrect. If Brad adopts the generosity policy and Chad adopts the caution policy, then Ann will receive all of the recommendation information from Brad. However,

because Chad does not have a recommendation for Sam's trustworthiness since he believes that his only source, Eve, is dishonest, Ann will not receive any information from Chad. Because Ann does not have any information leading her not to trust Eve, she will believe the information from Brad, leading again to the incorrect result. If both Chad and Brad adopt the caution policy, then Ann will receive a recommendation that Sam is trustworthy from Brad and no information from Chad, so Ann will accept Brad's recommendation and again incorrectly believe that Sam is highly trustworthy. Therefore, only if both Brad and Chad adopt the generosity policy towards Ann will Ann be able to avoid being tricked by Eve through the application of the RTI algorithm.

In general, the addition of privacy policies to the trust model enables each agent to adopt a different kind of behavior towards each other agent based on the level of trust it has for the other agent. For example, if one agent has reason to believe that another agent asking for information is dishonest, it can choose to adopt the "non-cooperation policy" which in the RTI algorithm means that no recommendation is returned to the requesting agent. If the agent does not know anything about the requesting agent, it might adopt the "caution policy" as a default, giving the requesting agent only the minimum information needed to make the trust network work and hiding the path information of the other agents between it and the target agent. If the agent believes that the requesting agent is trustworthy, perhaps because another trustworthy agent has vouched for it, the agent can adopt the "generosity policy", in which case it returns the recommendation and the path information that it receives from all of the agents between it and the target agent. Then, the requesting agent can interpret all of the information that it receives in terms of the original RTI algorithm and calculate the inferred trust value for the target agent as described in [5].

Now, we consider a slightly more complicated situation where Eve knows more than two agents and where two of the agents she knows give the same trust recommendation for the target agent. The social network shown in Fig. 2 is composed of seven agents representing the Web users Ann, Brad, Chad, Diana, Eve, Kay, and Mary. Mary is a highly trustworthy service provider. There are eight relationships between the seven agents, forming three chains of trust links that connect Ann to Mary.

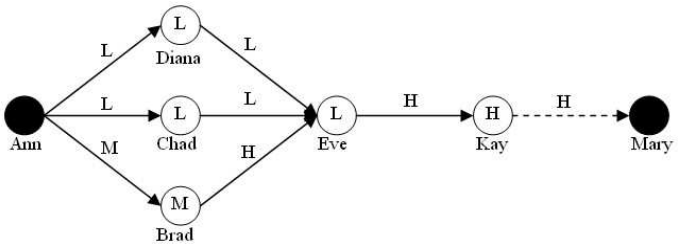

Fig. 2. The second scenario with three parallel trust chains

In this example, because Kay is the final user in all of the trust chains, Kay must adopt the generosity policy to Eve. Furthermore, if Eve adopts either the caution policy or the generosity policy to Diana, Chad and Brad, the final result will be the same. Therefore, we only need to consider the privacy policies adopted by Diana, Chad, and Brad to Ann. In here, we just consider the following two situations to illustrate the impact of the privacy policies to RTI algorithm.

If Brad and Chad adopt the generosity policy and Diana adopts the caution policy, then Ann will receive the single recommendation of "low trust" for Eve from Chad, and the single recommendation of "high trust" for Eve from Brad, but the recommendation from Diana for Mary will have the aggregated reliability value of $0.7 \times 0.9 = 0.63$ with no information about who gave the recommendation. Therefore, Ann will calculate the trust value for Eve to be "high trust" because Ann trusts Brad more than Chad for a single recommendation, i.e. Ann calculates the reliability of the second trust chain using H (0.9) replacing L (0.7) for the trust value of Chad to Eve. The final result for the reliability of the recommendation that Mary is highly trustworthy is $1 - (1 - 0.7 \times 0.63) \times (1 - 0.7 \times 0.9 \times 0.9) \times (1 - 0.8 \times 0.9 \times 0.9) = 0.91$.

If Brad adopts the generosity policy, and Chad and Diana both adopt the caution policy, then Ann will receive the single recommendation of "high trust" for Eve from Brad. The recommendations for Mary from Chad and Diana will not have any information about who gave the recommendations. Therefore, Ann will assign the trust value for Eve to be "high trust", but she will use the aggregate reliabilities for the recommendations of Mary's trustworthiness from Chad and Diana of $0.7 \times 0.9 = 0.63$. The final result for the reliability of the recommendation that Mary is highly trustworthy is $1 - (1 - 0.7 \times 0.63) \times (1 - 0.7 \times 0.63) \times (1 - 0.8 \times 0.9 \times 0.9) = 0.89$.

The analyses above show the effect of protecting privacy in the trust network. If the agents adopt the caution policy, information from some trust chains will be lost, and the accuracy of the inferred trust value will decrease. On the other hand, if an agent adopts the generosity policy, then it risks having its privacy information exposed. Our implementation of the RTI algorithm with privacy protection supports dynamic propagation of trust and privacy information in two ways. Whenever an agent receives trust recommendations from highly trusted agents for agents that it knows but does not trust, our implementation allows the agent to update the privacy policies accordingly. Alternatively, any time a person represented by an agent confirms that another agent is either trustworthy or dishonest, that person can manually assign an appropriate privacy policy. Furthermore, when an agent changes its trust level for another agent in either of these ways, it will send the new trust information to all of the agents to which it has adopted the generosity policy, resulting in a push style of trust information transfer. This push style information transfer will only occur between highly trusted agents adopting the generosity policy to each other. Each community of agents will adopt its own guidelines for balancing the risks of trusting a particular agent against the benefits of getting useful information from that agent using all of the trustworthy information that is available to it,

much the same way that humans interact in society. Our hope is that this will result in the establishment of reliable recommendation networks where a recommendation for a particular agent will be updated quickly among the highly trusted peers of the recommending agent. Because each agent knows that if it is dishonest, its malicious reputation will be rapidly spread through the peer-based connections of the network, we hypothesize that most agents will be motivated to stay honest and friendly. In this way, we propose that the privacy policies in the trust network will result in a dynamic equilibrium where most agents are honest and adopt generosity policy between each other, which forms a robust network of trust recommendation that rapidly exposes dishonest agents, keeping their numbers down. Then, a high accuracy of inferred trust can be maintained while simultaneously protecting privacy.

5 Future Work

We plan to continue our research along several directions. First, we will create a trust network that closely simulates real social networks by exhibiting characteristics such as small world behavior. Based on that, we will conduct simulation studies to analyze how the different trust metrics work. Second, we are investigating methods for integrating the trust-recommendation network with the EKOSS knowledge searching and matching system in order to share different quantities and qualities of knowledge with agents that have different trust values.

References

1. Kraines, S.B., Batres, B., Koyama, M., Wallace, D. R., and Komiyama, H., Internet-based collaboration for Integrated Environmental Assessment in Industrial Ecology - Part 1, Journal of Industrial Ecology 9(3) (2005) 31-50.
2. Kraines, S.B., Batres, R., Kemper, B., Koyama, M., and Wolowski, V. (2006) 'Internet-Based Integrated Environmental Assessment, Part II: Semantic Searching Based on Ontologies and Agent Systems for Knowledge Discovery', Journal of Industrial Ecology, Vol. 10, No. 4, pp.1-24.
3. EKOSS site, (http://www.ekoss.org) World Wide Web.
4. Kraines, S., Guo, W., Kemper, B., and Nakamura, Y., EKOSS: A Knowledge-User Centered Approach to Knowledge Sharing, Discovery, and Integration on the Semantic Web. Proc. of ISWC 06, LNCS 4273/2006, pp. 833-846
5. Guo, W., Kraines, S., Inferring Trust from Recommendations in Web-based Knowledge Sharing Systems. Adv. in Intel. Web, ASC43/2007, pp. 148-153.
6. Langheinrich, M., Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. Ubicomp 2001, LNCS 2201/2001, pp. 273-291.
7. Goecks, J., and Mynatt, E., Enabling Privacy Management in Ubiquitous Computing Environments through Trust and Reputation Systems. Proc. of the 2002 ACM Conference on Computer Supported Cooperative Work, New Orleans, LA, USA.
8. Golbeck, J., Hendler, J., Accuracy of Metrics for Inferring Trust and Reputation in Semantic Web-based Social Networks. LNCS, 3257/2004, pp. 116-131.