

Trust and Reputation Model in Peer-to-Peer Networks

Yao Wang, Julita Vassileva
University of Saskatchewan, Computer Science Department,
Saskatoon, SK, S7N 5A9, Canada
{yaw181, jiv}@cs.usask.ca

Abstract

It is important to enable peers to represent and update their trust in other peers in open networks for sharing files, and especially services. In this paper, we propose a Bayesian network-based trust model and a method for building reputation based on recommendations in peer-to-peer networks. Since trust is multi-faceted, peers need to develop differentiated trust in different aspects of other peers' capability. The peer's needs are different in different situations. Depending on the situation, a peer may need to consider its trust in a specific aspect of another peer's capability or in multiple aspects. Bayesian networks provide a flexible method to present differentiated trust and combine different aspects of trust. The evaluation of the model using a simulation shows that the system where peers communicate their experiences (recommendations) outperforms the system where peers do not share recommendations with each other and that a differentiated trust adds to the performance in terms of percentage of successful interactions.

1. Introduction

Peer-to-peer networks are networks in which peers cooperate to perform a critical function in a decentralized manner [6]. All peers are both consumers and providers of resources and can access each other directly without intermediary peers. Compared with a centralized system, a peer-to-peer (P2P) system provides an easy way to aggregate large amounts of resources residing on the edge of Internet or in ad-hoc networks with a low cost of system maintenance. P2P systems have attracted increasing attention from researchers recently, but they also bring up some problems. Since peers are heterogeneous, some peers might be benevolent in providing services. Some might be buggy or malicious and cannot provide services with the quality that they advertise. Since there is no centralized node to serve as an authority to monitor and punish the peers that behave badly, malicious peers have an incentive to provide poor quality services for their benefit because

they can get away. Some traditional security techniques, such as service providers requiring access authorization, or consumers requiring server authentication, are used as protection from known malicious peers. However, they cannot prevent from peers providing variable-quality service, or peers that are unknown. Mechanisms for trust and reputation can be used to help peers distinguish good from bad partners. This paper describes a trust and reputation mechanism that allows peers to discover partners who meet their individual requirements through individual experience and sharing experiences with other peers with similar preferences.

The rest of this paper is organized as follows: section 2 discusses the definitions of trust and reputation and their characteristics. Section 3 introduces our approach to developing a Bayesian network-based trust model and a method for building reputation based on recommendations. The experiment design and results are presented in Sections 4 and 5. Section 6 discusses related work on trust and reputation. In the last section, we present conclusions and directions for future work.

2. Trust and reputation

Trust and reputation mechanisms have been proposed for large open environments in e-commerce, peer-to-peer computing, recommender systems [4, 13, 14, 17, 18, 19]. However, there is no universal agreement on the definition of trust and reputation. In this paper, we adopt the following working definitions:

Trust – a peer's belief in another peer's capabilities, honesty and reliability *based on its own direct experiences*;

Reputation – a peer's belief in another peer's capabilities, honesty and reliability *based on recommendations* received from other peers. Reputation can be centralized, computed by a trusted third party, like a Better Business Bureau; or it can be decentralized, computed independently by each peer after asking other peers for recommendations.

Although trust and reputation are different in how they are developed, they are closely related. They are both used

to evaluate a peer's trustworthiness, so they also share some common characteristics [1, 8, and 12].

➤ *Context specific.*

Trust and reputation both depend on some context. For example, Mike trusts John as his doctor, but he does not trust John as a mechanic who can fix his car. So in the context of seeing a doctor, John is trustworthy. But in the context of fixing a car, John is untrustworthy.

➤ *Multi-faceted.*

Even in the same context, there is a need to develop differentiated trust in different aspects of the capability of a given peer. The same applies for reputation. For instance, a customer might evaluate a restaurant from several aspects, for example, the quality of food, the price, and the service. For each aspect, she develops a kind of trust. The overall trust depends on the combination of the trusts in each aspect. While the *context-specificity* of trust accentuates that trust in an identical peer can be different in different situations, the characteristic, *multi-faceted*, emphasizes that trust has multiple aspects, which can play a role in deciding whether a peer is trustworthy to interact with.

➤ *Dynamic.*

Trust and reputation increase or decrease with further experience (direct interaction). They also decay with time.

3. Bayesian network-based trust model

3.1 Trust and reputation mechanism

In our model a peer builds two kinds of trust in another peer, say peer A and peer B respectively. The first one is the trust that peer A has in peer B's *capability in providing services*. The other is the trust that peer A has in peer B's *reliability in providing recommendations* about other peers. Here the reliability includes two aspects:

- *Truthfulness* – whether peer B is truthful in telling its information
- *Similarity* – whether peer B is similar to peer A in preferences and ways of judging issues.

Reliability = Truthfulness ? Similarity, i.e. a peer B's reliability as a referee depends on both being truthful and similar in its preferences to the peer requesting the recommendation. Since peers are heterogeneous, they may have different preferences and judge issues by different criteria. For example, some peers may consider a movie provider good because it provides movies with high quality, while others may consider the movie provider bad because the speed of download from it is very slow. If two peers A and B are similar in their evaluation criteria, peer A can trust peer B's recommendations, if it knows that peer B is truthful. However, if the peers have different evaluation criteria, peer A cannot trust peer B's recommendations even when peer B tells the truth.

It is important for a peer to develop trust in other peers as references in a decentralized system, since when a peer

is not sure about the trustworthiness of a service provider, it can ask for recommendations only those few peers that it trusts most instead of asking a large number of peers, which not only helps the peer get more reliable recommendations, but also saves time and communication costs.

We will use a peer-to-peer file sharing application as an example in the discussion, however the method is general and can be applied to other applications, like web-services, e-commerce, recommender systems or peer-to-peer distributed computing.

In the area of file sharing in peer-to-peer networks, all the peers are both providers and users of shared files. Each peer plays two roles, the role of file provider offering files to other peers and the role of user searching and downloading files provided by other peers. In order to distinguish the two roles of each peer, in the rest of paper, when a peer acts as a file provider, we call it file provider; otherwise, we call it simply peer. Peers will develop two kinds of trust, the trust in the file providers' capability (in providing files) and the trust in the other peers' reliability in making recommendations. We assume all the peers are truthful in telling their evaluations. However, the peers may have different ways of evaluating other peers' performance, which reflect different user preferences.

A search request in file sharing peer-to-peer applications usually results in a long list of providers for an identical file. If a peer happens to select a provider of files with bad quality or slow download speed, the peer will waste time and effort, which may lead to user frustration and abandoning the system. In order to solve the problem, we use the mechanism of trust and reputation as shown in Figure 1. Once a peer receives a list of file providers for a given search, it can arrange the list according to its trust in these file providers. Then the peer chooses one of the file providers on top of the list. If the file provider is trustworthy according to the peer's previous experiences, the peer will interact with the file provider (download files). If the file provider is not trustworthy, the peer will select another file provider to interact with. If the peer is not sure about the trustworthiness of the file provider, for example, the peer has no interactions or only a few interactions with the file provider, it can ask other peers to make recommendations for it. How the peer uses the reputation and its own trust to make a decision with which file provider to interact is an open question. Some peers may prefer to trust their own experience and rely on their trust even if they had very few interactions with the service provider. Others may be more cautious and rely on the reputation of the service provider. After each interaction, the peer updates its trust in the file provider according to its evaluation of the interaction. If the interaction is satisfying, it will increase its trust in the file provider; if the interaction is not satisfying, it will decrease its trust in the file provider. If the decision of interaction is based on other peers' recommendations, the peer will also update its trust

in each of the peers that give recommendations (we call these peers “referees”). If the referee’s recommendation is consistent with the peer’s evaluation of the interaction, the peer will increase its trust in the referee; otherwise, it will decrease its trust.

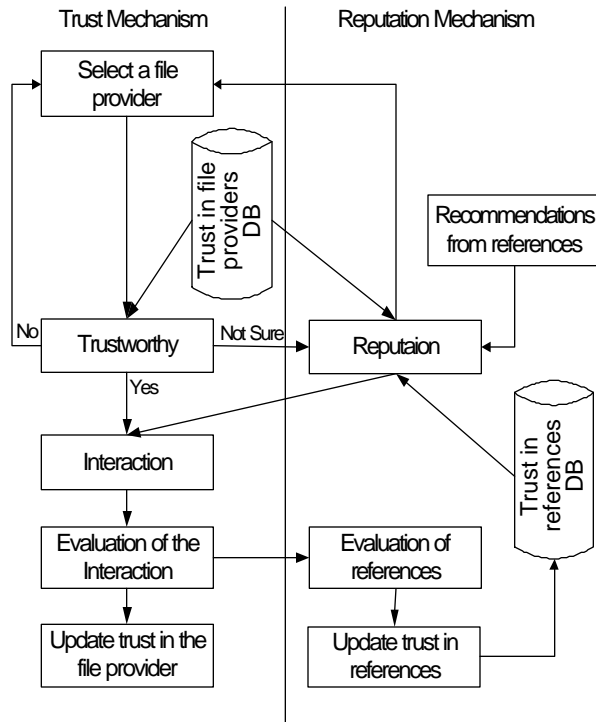


Figure 1. Functionality of the trust and reputation mechanism on board of the peer

3.2 Trust in a file provider’s capability

In a peer-to-peer network, file providers’ capabilities are not uniform. For example, some file providers (FP) may be connecting through a high-speed network, while others connect through a slow modem. Some file providers might like music, so they share a lot of music files. Some may be interested in movies and share more movies. Some may be very picky about file quality, so they only keep and share files with high quality. Therefore, the file provider’s capability can be presented in various aspects, such as the download speed, file quality and file type.




Figure 2. A Bayesian network model

The peer’s needs are also different in different situations. Sometimes, it may want to know the file provider’s overall capability. Sometimes it may only be interested in the file provider’s capability in some particular aspect. For instance, a peer wants to download a music file from a file provider. At this time, knowing the file provider’s capability in providing music files is more valuable for the peer than knowing the file provider’s capability in providing movies.

Peers also need to develop differentiated trust in the file providers’ capabilities. For example, the peer who wants to download a music file from the file provider cares about whether the file provider is able to provide the music file with good quality at a fast speed, which involves the file provider’s capabilities in two aspects, quality and speed.

How does the peer combine its two separated trust representations together, the trust in the file provider’s capability in providing music files with good quality and the trust in the file provider’s capability in providing a fast download speed, in order to decide whether the file provider is trustworthy or not?

A Bayesian network provides a flexible method to solve the problem. It is a relationship network that uses statistic methods to represent probability relationships between different elements [10]. We use a naïve Bayesian network to represent the trust of a peer in a file provider. Every peer develops a naïve Bayesian network for each file provider that it has interacted with. Each Bayesian network (see Figure 2) has a root node *T* that represents the peer’s trust in the file provider’s capability in providing files. It is the percentage of interactions that are satisfying. The leaf nodes under the root node represent the file provider’s capability in different aspects. The node, denoted by *FT*, represents the set of file types. Suppose it includes five values, “*Music*”, “*Movie*”, “*Document*”, “*Image*” and “*Software*”. The node “*DS*” denotes the set of download speeds. It has three values, “*Fast*”, “*Medium*” and “*Slow*”, each of which covers a range of download speeds. The node “*FQ*” denotes the set of file qualities. It also has three values, “*High*”, “*Medium*” and “*Low*”.

Here we only take three aspects of trust into account. More relevant aspects can be added in the Bayesian network later to account for user preferences with respect to service.

According to a Bayesian network, a peer can infer the trustworthiness of a file provider in different conditions, such as the trustworthiness of the file provider in providing music files, the trustworthiness of the file provider in providing files with high quality, the trustworthiness of the file provider in providing music files with high quality. The condition can be any combination of the aspects. The method will save peers effort in building different trusts separately, or developing new trust when conditions change.

3.3 Evaluating interactions and updating trust in file providers

After each interaction, peers make an evaluation of it. Peers might have different criteria to judge an interaction. Some peers might be very picky. Some might be generous. So they might have different evaluations of an identical interaction. The overall evaluation of an interaction is a combination of evaluations of each aspect related to the interaction, such as download speeds, file quality. How to combine evaluations of each aspect depends on each peer's preference. For example, some peers more care about the download speed. Some more care about the quality of downloaded files. Some may equally care about both of them.

The result of the overall evaluation, "*the interaction is satisfying*" or "*not satisfying*", is used to update the peer's trust in the file provider involved. The update is implemented by adding the new experience into the peer's corresponding Bayesian network. The details are shown in [16].

3.4 Handling recommendations

When a peer is not sure about the trustworthiness of a file provider, it can ask other peers for recommendations. The recommendation requests can vary according to the peer's needs. For example, if the peer is going to download a movie, it may care about the movie's quality. Another peer may care about the download speed. So the request can be "Does the file provider have movies with good quality?" If the peer cares both about the quality and the download speed, the request will be something like "Does the file provider offer files with good quality and fast download speed?" When other peers receive these requests, they will check their trust representations, i.e. their Bayesian networks, to see if they can answer such questions. If a peer has downloaded movies from the file provider before, it will answer the first question with its trust in the file provider under the condition that the file

provider provides files with good quality and the second question with its trust under the condition that the file provider provides files with good quality and fast download speed according to its Bayesian network.

The peer might receive several such recommendations at the same time from trustworthy, untrustworthy acquaintances, or strangers. If the references are untrustworthy, the peer can discard their recommendations immediately. Then the peer needs to combine the recommendations from trustworthy references and from unknown references to get the total recommendation for the file provider. Peers may value the importance of the recommendation from trustworthy references and from unknown references differently. Since peers often have different preferences and points of view, the peer's trustworthy acquaintances are those peers that share similar preferences and viewpoints with the peer most of time. The peer should weight the recommendations from its trustworthy acquaintances higher than the recommendations from strangers. Given a threshold q , if the total recommendation value is greater than q , the peer will interact with the file provider; otherwise, not.

If the peer interacts with the file provider, it will not only update its trust in the file provider, i.e. its corresponding Bayesian network, but also its trust in the referee-peers that provide recommendations by the following reinforcement learning formula:

$$tr_{ij}^n = \mathbf{a} * tr_{ij}^o + (1 - \mathbf{a}) * e_a \quad (1)$$

tr_{ij}^n denotes the new trust value that the i^{th} peer has in the j^{th} referee after the update; tr_{ij}^o denotes the old trust value. \mathbf{a} is the learning rate – a real number in the interval [0,1]. e_a is the new evidence value, which can be -1 or 1. If the value of recommendation is greater than q and the interaction with the file provider afterwards is successful, e_a is equal to 1. If there is a mismatch between the recommendation and the actual experience with the file provider, the evidence is negative, so e_a is -1.

Another way to find if a peer is reliable in making recommendations is the comparison between two peers' Bayesian networks relevant to an identical file provider. When peers are idle, they can "gossip" with each other periodically, exchange and compare their Bayesian networks. This can help them find other peers who share similar preferences more accurately and faster. After each comparison, the peers will update their trusts in each other according the formula:

$$tr_{ij}^n = \mathbf{b} * tr_{ij}^o + (1 - \mathbf{b}) * e_b \quad (2)$$

The result of the comparison e_b is a number in the interval [-1, 1]. \mathbf{b} is the learning rate – a real number in the interval [0,1], which follows the constraint $\mathbf{b} > \mathbf{a}$. This is because the Bayesian network collectively reflects a

peer's preferences and viewpoints based on all its past interactions with a specific file provider. Comparing the two peers' Bayesian networks is tantamount to comparing all the past interactions of the two peers. The evidence e_a in formula (1) is only based on one interaction. The evidence e_b should affect the peer's trust in another peer more than e_a .

How do the peers compare their Bayesian networks and how is e_b computed? First, we assume all peers have the same structure of Bayesian networks. We only compare the values in their Bayesian networks. Suppose peer 1 will compare its Bayesian network with the corresponding Bayesian network of peer 2. Peer 1 gets the degree of similarity between the two Bayesian networks by computing the similarity of each pair of nodes (T, DS, FQ and FT), according to the similarity measure based on Clark's distance [7], and then combining the similarity results of each pair of nodes with different weight in order to take into account peers' preferences. So peers with similar preferences, such as the importance of file type, quality, and download speed, will weight each other's opinions higher.

In the above discussion, we assume all the peers are truthful in making recommendations. In the situation that peers are not truthful, our method is still suitable. Since a file provider's reputation is built on a collection of recommendations, even if a few peers lie, it will not influence the overall reputation of the file provider. If a peer does lie to another peer, for example, peer A lies to peer B, peer B's trust in peer A as a referee will decrease quickly because peer A's recommendation does not match peer B's evaluation of the involved interaction.

4. Experiments

In order to evaluate this approach, we developed a simulation of a file sharing system in a peer-to-peer network. The system is developed on the JADE 2.5. For the sake of simplicity, each node in our system plays only one role at a time, either the role of file provider or the role of a peer. At the beginning every peer knows only peers directly connected with it and a few file providers.

Every peer has an interest vector. The interest vector is composed of five elements: *music*, *movie*, *image*, *document* and *software*. The value of each element indicates the strength of the peer's interests in the corresponding file type. The files the peer wants to download are generated based on its interest vector. Every peer keeps two lists. One is the peer list that records all the other peers that the peer has interacted with and its trust values in these peers. The other is the file provider list that records the known file providers and the corresponding Bayesian networks representing the peer's trusts in these file providers. Each file provider has a capability vector

showing its capabilities in different aspects, i.e. providing files with different types, qualities and download speeds.

Our experiments involve 10 different file providers and 40 peers. Peers will gossip with other peers periodically (after every 5 interactions) to exchange their Bayesian networks. The total number of interactions is 1000. We run each configuration for 10 times and use means for the evaluation criteria.

5. Results

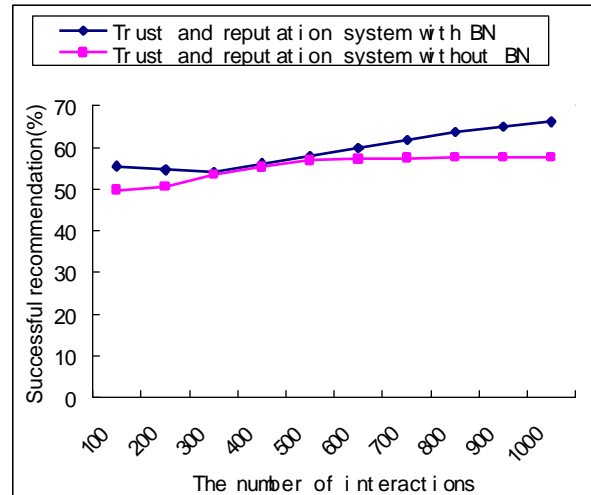


Figure 3. Trust and reputation system with BN vs. trust and reputation system without BN

The goal of the first experiment is to see if a Bayesian network-based trust model helps peers to select file providers that match better their preferences. Therefore we measure the system performance in terms of percentage of successful recommendations. A successful recommendation is defined as a positive recommendation about a file provider such that, after receiving it and interacting with the file provider, the peer is satisfied with the interaction. The percentage of successful recommendation is the number of successful recommendations divided by the number of positive recommendations because if a peer gets a negative recommendation for a file provider, it will not interact with the file provider. So we are looking at the proportion of satisfactory performance over unsatisfactory performance after positive recommendation.

We compare the performance of a system consisting of peers with Bayesian network-based trust models and a system consisting of peers without Bayesian networks (BN) trust model. These peers represent general trust only, which is not differentiated into different aspects. So, we have two configurations in this experiment:

- *Trust and reputation system with BN*: the system consists of peers with Bayesian networks-based trust

models that exchange recommendations with each other;

- *Trust and reputation system without BN*: the system consists of peers that exchange recommendations, but do not model differentiated trust in file providers;

Figure 3 shows that the system using Bayesian networks performs slightly better than the system with general trust in terms of the percentage of successful recommendations.




Figure 4. The comparison of four systems

The goal of the second experiment is to see if exchanging recommendation values with other peers helps peers to achieve better performance defined as the percentage of successful interactions with file provider, which is the number of successful interactions over the total number of interactions. For the reason, we compare four configurations:

- Trust and reputation system with BN;
- Trust and reputation system without BN;
- Trust system with BN: the system consists of peers with Bayesian networks-based trust models, which do not exchange recommendations with each other;
- Trust system without BN: the system consists of peers that have no differentiated trust models and do not exchange recommendations with each other.

Figure 4 shows that the two systems, where peers share information with each other, outperform the systems, where peers do not share information. The trust system using Bayesian networks is slightly better than the trust system without using Bayesian networks. There is an anomaly in the case when peers do not share recommendations, since in the end of the curve, the system without BN perform better than the system with BN. This could be explained with an imprecise BN due to insufficient experience.

In some sense, a peer's Bayesian network can be viewed as the model of a specified file provider from the

peer's personal perspective. In our experiments, we use a very simple naive Bayesian network, which cannot represent complex relationships. In the real file-sharing system, the model of file providers might be more complex and required using a more complex Bayesian network. Our Bayesian network only involves three factors. In future, we will build a more complex Bayesian network and add more aspects into it to see how the system works.




Figure 5. Tendency in trust and reputation systems with BN

The goal of the third experiment is to see the tendencies about successful recommendations and successful interactions in trust and reputation systems with BN. The number of interactions this time is 2000. We measured the following three parameters:

- *R-Succ/Recom*, the percentage of successful recommendations, which is the number of successful recommendations divided by the number of positive recommendations.
- *R-Succ/Succ*, the percentage of successful interactions based on recommendations, which is the number of successful interactions based on recommendations over the number of all successful interactions.
- *Succ/Inter*, the percentage of successful interactions in all interactions, which is the number of successful interactions divided by the total number of interactions.

Figure 5 shows that *R-Succ/Recom* and *Succ/Inter* tend to be stable with the increase of the number of interactions, which indicates that the percentages of successful recommendations and successful interactions are going to reach their maximal values determined by the capabilities of file providers. *R-Succ/Succ* tends to decrease with the increase of number of interactions, which suggests that peers need less and less recommendations when they have enough experiences with file providers.

6. Discussion and related work

How many Bayesian networks can a peer afford to maintain to represent its trust in other peers in the networks? It depends on the size of the network and the likelihood that peers have repeated interactions. Resnick [11] empirically shows that 89.0% of all seller-buyer pairs in eBay conducted just one transaction during a five-month period and 98.9% conducted no more than four. This situation often happens in a very large network or in large e-commerce sites. Since there are a large number of sellers and buyers, the chance that a buyer meets the same seller is small. But if the kind of goods being transacted is only interesting to a small group of people, for example, collectors of ancient coins, the interactions about this kind of goods happen almost exclusively in a small group. So the probability that sellers and buyers have repeated interactions will be high, and they will be able to build trust in each other by our method. Keeping Bayesian networks trust models of a relatively small group of peers will not be too expensive.

Our approach is useful in situations where two peers can repeatedly interact with each other. In a small-size network, there is no doubt that our approach is applicable. For a large network, our approach is still suitable under the condition that the small-world phenomenon happens. The small-world phenomenon was first discovered in the 1960ies by social scientists. Milgram's experiment showed that people in the U.S. are connected by a short (average length of 6) chain of intermediate acquaintances. Other studies have shown that people tend to interact with other people in their small world more frequently than with people outside. The phenomenon also happens in peer-to-peer networks. Jovanovic's work [5] proves that the small-world phenomenon occurs in Gnutella. It means that peers are inclined to get files from other peers from a small sub-community. This small sub-community often consists of peers that have similar preferences and viewpoints.

The trust and reputation mechanism can not only help peers find trustworthy file providers, but also automatically balance the workload of file providers. For example, if a lot of peers download files from a file provider at the same time, the download speed that each peer gets will decrease, which leads to the peers' bad evaluations of the interactions with the file provider. As a result, the peers' trust in the file provider will reduce, which causes the decrease of the reputation of the file provider and the decrease of peers that download files from it. Accordingly, the workload of the file provider will be reduced and will be shifted to some other file providers.

There is a lot of research on trust and reputation. Here we just mention some works that are most related to our approach.

Abdul-Rahman and Hailes [1] capture the most important characteristics of trust and reputation and propose the general structure for developing trust and reputation in a distributed system. Most of the later works in the area follow their ideas, but in different application domain, such as [2, 3, 4, 7]. Sabater and Sierra's work [12] extends the notion of trust and reputation into social and ontological dimensions. Social dimension means that the reputation of the group that an individual belongs to also influences the reputation of the individual. Ontological dimension means that the reputation of a peer is compositional. The overall reputation is obtained as a result of the combination of the peer's reputation in each aspect.

Our approach integrates these two previous works [1, 12], and applies them to file sharing system in peer-to-peer networks. Another difference between our work and Sabater and Sierra's work is that we use Bayesian networks to represent trust in different aspects, other than the structure of ontology. Another difference is that we do not treat the differentiated trusts as compositional. Usually the relationship between different aspects of a peer is not just compositional, but complex and correlative. Our approach provides an easy way to present a complex and correlative relationship. Our approach is also flexible in inferring the trust of a peer for different needs. For example, sometimes we care about the overall trust. Sometimes we only need to know the trust in some specific aspect. This bears parallel with work on distributed user modeling and purpose-based user modeling [9, 15].

Cornelli's work [4] is also in the area of file sharing in peer-to-peer networks. However, it concentrates on how to prevent attacks on the reputation system and does not discuss how peers model and compute trust and reputation.

7. Conclusions

Enabling peers to develop trust and reputation among themselves is important in a peer-to-peer system where resources (either computational, or files) of different quality are offered. It will become increasingly important in systems for peer-to-peer computation, where trust and reputation mechanisms can provide a way for protection of unreliable, buggy, infected or malicious peers. In this paper, we propose a Bayesian network-based trust model and a method for building reputation based on recommendations in peer-to-peer networks. Bayesian networks provide a flexible method to present the differentiated trust and combine different aspects of trust. In order to evaluate our approach, we developed a simulation of a file sharing system in a peer-to-peer network. Our experiments show that the system where peers communicate their experiences (recommendations) outperforms the system where peers do not communicate with each other

and that a differentiated trust adds to the performance in terms of percentage of successful interactions.

Future work includes adding more aspects in the Bayesian networks, trying to find the key parameters that influence the system performance, and testing the system under other performance measures, for example, how fast a peer can locate a trustworthy service provider and how fast the workload of file providers can be balanced. Applying this approach to peer-to-peer systems for computational services is particularly promising.

Acknowledgements

This research is funded by the NSERC Individual Discovery Grant of the second author.

References

- [1] Abdul-Rahman A. and Hailes S. "Supporting trust in virtual communities". In Proceedings of the Hawai'i International Conference on System Sciences, Maui, Hawaii, Jan 4-7 2000.
- [2] Azzedin F. and Maheswaran M. "Evolving and Managing Trust in Grid Computing Systems". IEEE Canadian Conference on Electrical & Computer Engineering (CCECE '02), May 2002.
- [3] Carter J., Bitting E. and Ghorbani A. "Reputation Formalization for an Information-Sharing Multi-Peer System", 515-534. Computational Intelligence, Volume 18, Number 4, November 2002.
- [4] Cornelli F. and Damiani E. "Implementing a Reputation-Aware Gnutella Servent". In Proceedings of the International Workshop on Peer-to-Peer Computing, Pisa, Italy, May 24, 2002.
- [5] Jovanovic M. "Modeling Large-scale Peer-to-Peer Networks and a Case study of Gnutella", University of Cincinnati, master thesis, April 2001.
- [6] Milojicic D. S., Kalogeraki V. and Lukose R. "Peer-to-Peer Computing", Tech Report: HPL-2002-57, available on line at: <http://www.hpl.hp.com/techreports/2002/HPL-2002-57.pdf>
- [7] Montaner M. and López B. "Opinion based filtering through trust". In Proceedings of the 6th International Workshop on Cooperative Information Peers (CIA'02), Madrid (Spain), September 18-20 2002.
- [8] Mui L., Halberstadt A. and Mohtashemi M. "Notions of reputation in multi-peers systems: A review". In Proceedings of Autonomous Peers & Multipeer Systems (AAMAS'02), 280-287, Bologna, Italy, 2002..
- [9] Niu X., McCalla G., Vassileva J. (to appear) "Purpose-based User Modelling in a Multi-peer Portfolio Management System". In Proceedings of User Modeling UM03, Johnstown, PA, June 22-26, 2003.
- [10] Heckerman, D. "A Tutorial on Learning with Bayesian Networks", Microsoft Research report MSR-TR-95-06, 1995.
- [11] Resnick P. and Zeckhauser R. "Trust among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System". NBER Workshop on Empirical Studies of Electronic Commerce, 2000.
- [12] Sabater J. and Sierra C. "Regret: a reputation model for gregarious societies". In 4th Workshop on Deception, Fraud and Trust in Peer Societies, 2001.
- [13] Schafer B.J. Konstan A, J, and Riedl J. "Recommender Systems in E-Commerce". ACM Conference on Electronic Commerce (EC-99), November 3-5, 1999, Denver, CO.
- [14] Vassileva J., Breban S. and Horsch M. "Peer Reasoning Mechanism for Long-Term Coalitions Based on Decision Making and Trust". Computational Intelligence, Vol. 18, no. 4, 2002.
- [15] Vassileva J., McCalla G. and Greer J. (accepted 17 October 2001) "Multi-Peer Multi-User Modeling", to appear in User Modeling and User-Adapted Interaction.
- [16] Wang, Y., Vassileva, J. "Bayesian Network Trust Model in Peer-to-Peer Networks" (to appear). In Proceedings of Second International Workshop Peers and Peer-to-Peer Computing, July 14, 2003. Melbourne, Australia.
- [17] Yu B. and Singh P. M. "A social mechanism of reputation management in electronic communities". In Proceedings of Fourth International Workshop on Cooperative Information Peers, 154-165, 2000.
- [18] Yu B. and Singh P. M. "An Evidential Model of Distributed Reputation Management". In Autonomous Peers & Multipeer Systems (AAMAS'02), 294-301, Bologna, Italy, 2002.
- [19] Zacharia G. Moukas A. and Maes P. "Collaborative Reputation Mechanisms in Electronic Marketplaces" In 32nd Annual Hawaii International Conference on System Science (HICSS-32), 1999.