

Bayesian Network-Based Trust 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

In this paper, we first review trust and reputation mechanisms from different perspectives. Then 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.

Keywords

Peer-to-peer networks, trust, reputation, Bayesian networks

1. Introduction

Peer-to-peer networks are networks in which all peers cooperate with each other to perform a critical function in a decentralized manner [10]. All peers are both consumers and providers of resources and can access each other directly without intermediary agents. 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 and cannot provide services as they advertise. Some might be malicious by providing bad services. Since there is no centralized node to serve as an authority to supervise peers' behaviors and punish peers that behave badly, malicious peers have an incentive to harm other peers to get more 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 reviews trust and reputation mechanisms from different perspectives. 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 Section 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, distributed computing, recommender systems. Agents are often used to manage and reason about trust and reputation. However, there is no universal agreement on the definition of trust and reputation in the multi-agent system research community.

2.1 Definitions and Characteristics

In this paper, we adopt the following working definitions:

Trust – an agent's belief in another agent's capabilities, honesty and reliability based on its own direct experiences;

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

Although trust and reputation are different in how they are developed, they are closely related. They are both used to evaluate an agent's trustworthiness, so they also share some common characteristics [1, 11, and 16].

- *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 agent. 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 agent can be different in different situations, the characteristic, *multi-faceted*, emphasizes that trust has multiple aspects, which can play a role in deciding whether an agent is trustworthy to interact with.

- *Dynamic.*

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

Despite different contexts, trust can be broadly categorized by the relationships between the two involved agents [7].

- *Trust between a user and her agent(s)*

When a (personal) agent represents a human user, there could be cases that the agent does not act as its user expects [18]. How much a user trusts her agent determines how she delegates her tasks to the agent.

- *Trust in service provider.*

It measures whether a service provider can provide trustworthy services. The quality of service is the main concern in this case.

- *Trust in references*

References refer to the agents that make recommendations or share their trust values. It measures whether an agent can provide reliable recommendations. It emphasizes the similarity in preferences and ways of judging issues between two agents. The more similar the two agents, the more they trust each other in making recommendations.

It is important for an agent to develop trust in other agents as references in a decentralized system, since when an agent is not sure about the trustworthiness of a service provider, it can ask for recommendations only those few agents that it trusts most instead of asking a large number of agents, which not only helps the agent get more reliable recommendations, but also saves time and communication costs.

- *Trust in groups*

An agent can trust a group of other agents [6, 8, 12, and 19]. By modeling trust in different groups, an agent can decide to join a group that can bring it most benefit [19]. Hales [8] points that group reputation can be a powerful mechanism for the promotion of beneficent norms under the right

conditions. This kind of trust is also useful in helping an agent judge another agent according to its trust in the group that the other agent belongs to.

2.2 Centralized vs. Decentralized

Trust and reputation mechanisms have been implemented in many systems adopting either a centralized structure or a decentralized structure. Accordingly, the trust and reputation mechanisms used in the two kinds of systems are also different.

In centralized systems, such as in eBay and onSale, which are mainly seen in the area of e-commerce, the trust and reputation mechanisms are relatively simple and have some common characteristics.

- A centralized node acts as the system manager responsible for collecting ratings from both sides involved in an interaction.
- Agents' reputations are public and global. The reputation of an agent is visible to all the other agents.
- Agents' reputations are built by the system. There is no explicit trust model between agents.
- Less communication is required between agents. An agent only communicates with the centralized node to find out the other agents' reputations.

Despite the simplicity of the centralized reputation mechanisms, empirical results show these systems do encourage transactions between sellers and buyers. But there are some problems. One problem is that agents are usually reluctant to give negative ratings because they can see each other's ratings and are afraid of revenges [15]. Another problem is that if an agent has a bad reputation, it can discard its old identity, choose a new one and start as a beginner, getting rid of its poor reputation. The third problem is that agents can increase their reputations artificially by creating fake identities and having them to give themselves high ratings [24].

The trust and reputation mechanisms used in decentralized systems, for example, peer-to-peer networks, are more complex than those applied in centralized systems. They have the following characteristics [2, 3, and 5]:

- There is no centralized system manager to govern trust and reputation.
- Subjective trust is explicitly developed by each agent. Each agent is responsible for developing its own trust in other agents based on their direct interactions.
- No global or public reputation exists. If agent A wants to know agent B's reputation, it has to proactively ask other agents for their evaluations of B, then synthesize the ratings together to compute agent B's reputation. The reputation of agent B developed by A is personalized because agent A can choose which agents it will ask for evaluations of B, its trustworthy friends or all known agents. Agent A can also decide how to combine the collected evaluations together to

get agent B's reputation. For example, it can only combine the evaluations coming from trusted agents. Or it can weight differently the evaluations from trusted agents, unknown agents and even untrustworthy agents when it combines them together.

- A lot of communication is required between agents to exchange their evaluations.

In decentralized systems, agent A can get agent B's reputation based on its own knowledge of the truthfulness of agents that make recommendations for agent B. So it is difficult for agent B to increase its reputation artificially. Since only agent A can see the recommendations, the references can express their feelings truthfully, not worried about potential revenges. But the tradeoff is that agents have to conduct a lot of communication and computation.

2.3 Application Areas

Trust and reputation was first used in e-commerce systems to encourage transactions between strangers. The use of trust and reputation has extended to other areas, including peer-to-peer networks, which involve a lot of uncertainty about the reliability of both sides of interactions.

- E-commerce.

eBay and *onSale* are the most famous e-commerce systems using reputation management mechanisms to help people find trustworthy partners to interact with [17]. They accumulate feedback from both sides of each interaction and publish buyers' and sellers' reputations. Users with bad reputations will be punished because no one would like to interact with them further. Studies on eBay have shown that its reputation system does work and encourages interactions. Another application example is that trust can be used to help forming long-term coalitions consisting of customers and vendors who have compatible preferences and interests [19]. The mechanism for forming long-term coalitions suggests three strategies to help customers and vendors decide which coalition to join: an individually oriented strategy and two socially oriented strategies. The individually oriented strategy is that a customer or vendor prefers to be in the coalition with the vendor or customer who she trusts most. Socially oriented strategies refer to the strategies that a customer or vendor prefers to join the coalitions that she trusts most. The trust in a coalition could be a function of the cumulative trust for each member in the coalition, or a function of the number of trustworthy partners in the coalition.

- Distributed Computing

In distributed computing, security is one major concern in resource sharing. Azzedin and Maheswaran [3] propose a behavior trust model in the grid computing to help resource providers and consumers interact with each other more safely. The behavior trust is built on past experiences and is

the result of combining together direct trust based on direct experiences and reputation (objective trust) based on recommendations.

- File Sharing P2P System

Cornelli [5] proposes a robust reputation mechanism in Gnutella to prevent some well-known security threats to reputation-based systems, for example, an attack by forging witnesses that give high ratings. In this approach, the reputation manager, a component of each Gnutella servant, will verify the existence of each witness by a direct connection.

- Information Filtering

Montaner and López [11] suggest an opinion-based information filtering method through trust. Agents build their initial trust according to the similarity between their opinions about some common items. When discovering a new item, they will ask their trusted friends to make recommendations for it, then combine their recommendations together and finally decide whether to recommend the new item to the user.

3. Bayesian Network-Based Trust Model

There is a lot of research on reputation [4, 21, 22, 23], but the study of trust has not received enough attention.

3.1 Trust Model

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

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

Reliability = Truthfulness ? Similarity, i.e. an agent B's reliability as a reference depends on both being truthful and similar in its preferences to the agent requesting the reference. Since agents are heterogeneous, they may have different preferences and judge issues by different criteria. For example, some agents 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 agents A and B are similar in their evaluation criteria, agent A can trust agent B's recommendations, if it knows that agent B is truthful. However, if the agents have different evaluation criteria, agent A cannot trust agent B's recommendations even when agent B tells the truth. In the implementation of such a system based on trust and reputation, some issues have to be considered.

- 1) How does an agent model its user? The user ultimately sets the criteria by which an agent evaluates other agents. Each

user has different preferences and ways of judging the quality of interaction. In order to act as its user wants, an agent has to keep learning its user's preferences and behaviors. If an agent fails to do as what its user expects, it will be useless.

- 2) How is an interaction to be evaluated? Trust is built on the agent's direct interactions with other agents. For each interaction, the agent's degree of satisfaction of the interaction will directly influence its trust in the other agent involved in the interaction. Usually, an interaction has multiple aspects and can be judged from different points of view.
- 3) How does an agent update its trust in another agent?
- 4) When will an agent ask for recommendations about another agent that it is going to interact with?
- 5) How does an agent combine together the recommendations for a given agent coming from different references? Since the recommendations might come from trusted agents, non-trusted agents or strangers, an agent has to decide how to deal with them.
- 6) How does an agent decide whether another agent is trustworthy to interact with or not, according to its direct experiences or reputation, or both?
- 7) How does an agent develop and update its trust in a reference that makes recommendations?
- 8) How many kinds of trust does an agent need to develop with another agent in a single context? In most situations, agents need to develop multiple trust relationships with each other in order to evaluate each other from different perspectives. For example, agent A might trust agent B in providing music files with good quality. But agent A might not trust agent B in offering movie files with the same quality as music files.
- 9) How does an agent decide whom to ask for recommendations? It can query all known agents, randomly selected few agents to query, or just its most trusted agents.

Our approach will deal with all the issues above except the first one, which is beyond our scope, although it is important in the case when the agent is acting directly on user's behalf. 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 agent. Agents will develop

two kinds of trust, the trust in file providers' competence (in providing files) and the trust in other agents' reliability in making recommendations. We assume all the agents are truthful in telling their evaluations. However, the agents may have different ways of evaluating other agents' performance, which reflects different user preferences.

3.2 Trust in a File Provider's Competence

In a peer-to-peer network, file providers' capabilities are not uniform. For example, some file providers 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 (see Figure 1).




Figure 1. A Bayesian Network Model

The agent'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, an agent 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 agent than knowing the file provider's capability in providing movies. Agents also need to develop differentiated trust in file providers' capabilities. For example, the agent 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 agent 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 values. Its theoretical foundation is the Bayes' rule [14].

$$p(h|e) = \frac{p(e|h) \cdot p(h)}{p(e)}$$

$p(h)$ is the prior probability of hypothesis h ; $p(e)$ is the prior probability of evidence e ; $p(h|e)$ is the probability of h given e ; $p(e|h)$ is the probability of e given h .

We will use a simple Bayesian network, a naïve Bayesian network, to represent the trust between an agent and a file provider, which is composed of a root node and several leaf nodes (see Figure 1).

Every agent develops a naïve Bayesian network for each file provider that it has interacted with. Each Bayesian network has a root node T , which has two values, “satisfying” and “unsatisfying”, denoted by 1 and 0, respectively. $p(T = 1)$ represents the value of agent’s overall trust in the file provider’s competence in providing files. It is the percentage of interactions that are satisfying and measured by the number of satisfying interactions, m , divided by the total number of interactions, n . $p(T = 0)$ is the percentage of not satisfying interactions.

$$p(T = 1) = \frac{m}{n} \quad (1)$$

$$p(T = 1) + p(T = 0) = 1$$

The leaf nodes under the root node represent the file provider’s capability in different aspects. Each leaf node is associated with a conditional probability table (CPT). The node, denoted by FT , represents the set of file types. Suppose it includes five values, “Music”, “Movie”, “Document”, “Image” and “Software”. Its CPT is showed in table 1. It includes two columns of values. Each column follows one constraint, which corresponds to one value of the root node. The sum of values of each column is 1.

Table 1. The CPT of Node FT

	T = 1	T = 0
Music	$p(FT = "Music" T = 1)$	$p(FT = "Music" T = 0)$
Movie	$p(FT = "Movie" T = 1)$	$p(FT = "Movie" T = 0)$
Document	$p(FT = "Docu" T = 1)$	$p(FT = "Docu" T = 0)$
Image	$p(FT = "Image" T = 1)$	$p(FT = "Image" T = 0)$
Software	$p(FT = "Soft" T = 1)$	$p(FT = "Soft" T = 0)$

$p(FT = "Music" | T = 1)$ is the conditional probability with the condition that an interaction is satisfying. It measures the probability that the file involved in an interaction is a music file, given the interaction is satisfying. It can be computed according to the following formula:

$$p(FT = "Music" | T = 1) = \frac{p(FT = "Music", T = 1)}{p(T = 1)}$$

$p(FT = "Music", T = 1)$ is the probability that interactions are satisfying and files involved are music files.

$$p(FT = "Music", T = 1) = \frac{m1}{n}$$

$m1$ is the number of satisfying interactions when files involved are music files.

$p(FT = "Music" | T = 0)$ denotes the probability that the files are music files, given that the interactions are not satisfying. The probabilities for other file types in Table 1 are computed in a similar way.

Node DS denotes the set of download speeds. It has three items, “Fast”, “Medium” and “Slow”, each of which covers a range of download speed.

Node FQ denotes the set of file qualities. It also has three items, “High”, “Medium” and “Low”. Its CPT is similar to the one in table 1.

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.

Once getting nodes’ CPTs in a Bayesian network, an agent can compute the probabilities that the corresponding file provider is trustworthy in different aspects by using Bayes’ rules, such as $p(T = 1 | FT = "Music")$ – the probability that the file provider is trustworthy in providing music files, $p(T = 1 | FQ = "High")$ – the probability that the file provider is trustworthy in providing files with high quality, $p(T = 1 | FT = "Music", FQ = "High")$ – the probability that the file provider is trustworthy in providing music files with high quality. Agents can set various conditions according to their needs. Each probability represents trust in an aspect of the file provider’s competence. With the Bayesian networks, agents can infer trust in the various aspects that they need from the corresponding probabilities. That will save agents much effort in building each trust separately, or developing new trust when conditions change. After each interaction, agents update their corresponding Bayesian networks.

3.3 Evaluation of an Interaction

Agents update their corresponding Bayesian networks after each interaction. If an interaction is satisfying, m and n are both increased by 1 using formula (1). If it is not satisfying, only n is increased by 1. Two main factors are considered when agents judge an interaction, the degree of their satisfaction with the download speed s_{ds} and the degree of their satisfaction with the quality of downloaded file s_{fq} . The overall degree of agents’ satisfaction with an interaction s is computed as the following:

$$s = w_{ds} * s_{ds} + w_{fq} * s_{fq}, \quad \text{where } w_{ds} + w_{fq} = 1 \quad (2)$$

w_{ds} and w_{fq} denote weights, which indicate the importance of download speed and the importance of file quality to a particular agent (depending on the user's preferences). Each agent has a satisfaction threshold s_t . If $s < s_t$, the interaction is unsatisfying; otherwise, it is satisfying.

3.4 Handling Recommendations

In file sharing peer-to-peer applications users find files by using the search function. In most situations, they get a long list of providers for an identical file. If a user happens to select an unsuitable provider, who provides files with bad quality or slow download speed, the user will waste time and effort. If this situation happens several times, the users will be frustrated. In order to solve the problem, we use the mechanism of trust and reputation. Once an agent 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 agent chooses the most trusted file providers in the top of the list to download files from. If the agent has no interactions or only a few interactions with the file provider, it can ask other agents to make recommendations for it. How the agent uses the reputation and its own trust to make a decision with which file provider to interact is an open question (question 6 in section 3). Some agents 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. The agent can send various recommendation requests according to its needs. For example, if the agent is going to download a movie, it may care about the movie's quality. Another agent may care about the speed. So the request can be "Does the file provider provide movies with good qualities?". If the agent cares both about the quality and the download speed, the request will be something like "Does the file provider have files with good quality at a fast download speed?". When other agents receive these requests, they will check their trust-representations, i.e. their Bayesian networks, to see if they can answer such questions. If an agent has downloaded movies from the file provider before, it will send recommendation that contains the value $p(T = 1 | FT = "Music", FQ = "High")$ to answer the first request or the value $p(T = 1 | FT = "Music", FQ = "High", DS = "Fast")$ to answer the second request.

The agent might receive several such recommendations at the same time, which may come from the trustworthy acquaintances, untrustworthy acquaintances, or strangers. If the references are untrustworthy, the agent can discard their recommendations immediately. Then the agent needs to combine the recommendations from trustworthy references and from unknown references to get the total recommendation for the file provider:

$$r_{ij} = w_t * \frac{\sum_{l=1}^k tr_{il} * t_{lj}}{\sum_{l=1}^k tr_{il}} + w_s * \frac{\sum_{z=1}^g t_{zj}}{g}, \text{ where } w_t + w_s = 1 \quad (3)$$

r_{ij} is the total recommendation value for the j^{th} file provider that the i^{th} agent gets. k and g are the number of trustworthy references and the number of unknown references, respectively. tr_{il} is the trust that the i^{th} user has in the l^{th} trustworthy reference. t_{lj} is the trust that the l^{th} trustworthy reference has in j^{th} file provider. t_{zj} is the trust that the z^{th} unknown reference has in j^{th} file provider. w_t and w_s are the weights to indicate how the user values the importance of the recommendation from trustworthy references and from unknown references. Since agents often have different preferences and points of view, the agent's trustworthy acquaintances are those agents that share similar preferences and viewpoints with the agent most of time. The agent should weight the recommendations from its trustworthy acquaintances higher than those recommendations from strangers. Given a threshold q , if the total recommendation value is greater than q , the agent will interact with the file provider; otherwise, not.

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

$$tr_{ij}^n = \mathbf{a} * tr_{ij}^o + (1 - \mathbf{a}) * e_a \quad (4)$$

tr_{ij}^n denotes the new trust value that the i^{th} agent has in the j^{th} reference after the update; tr_{ij}^o denotes the old trust value. \mathbf{a} is the learning rate c 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 satisfying, 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 an agent is reliable in making recommendations is the comparison between two agents' Bayesian networks relevant to an identical file provider. When agents are idle, they can "gossip" with each other periodically, exchange and compare their Bayesian networks. This can help them find other agents who share similar preferences more accurately and faster. After each comparison, the agents will update their trusts in each other according the formula:

$$tr_{ij}^n = \mathbf{b} * tr_{ij}^o + (1 - \mathbf{b}) * e_b \quad (5)$$

The result of the comparison e_b is a number in the interval $[-1, 1]$.

b is the learning rate – a real number in the interval $[0,1]$, which follows the constraint $b > a$. This is because the Bayesian network collectively reflects an agent’s preferences and viewpoints based on all its past interactions with a specific file provider. Comparing the two agents’ Bayesian networks is tantamount to comparing all the past interactions of the two agents. The evidence e_a in formula (4) is only based on one interaction. The evidence e_b should affect the agent’s trust in another agent more than e_a .

How do agents compare their Bayesian networks and how is e_b computed? First, we assume all agents have the same structure of Bayesian networks. We only compare the values in their Bayesian networks. Suppose agent 1 will compare its Bayesian network with the corresponding Bayesian network of agent 2. Agent 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 [12], and then combining the similarity results of each pair of nodes.

$$e_b = 1 - 2 * \sum_{i=1}^4 (w_{1_i} * c_i), \text{ where } w_{1_1} + w_{1_2} + w_{1_3} + w_{1_4} = 1 \quad (6)$$

$$c_1 = \frac{\sqrt{(v_{111} - v_{211})^2 + (v_{112} - v_{212})^2}}{\sqrt{(v_{111} + v_{211})^2 + (v_{112} + v_{212})^2}} \quad (7)$$

$$c_i = \frac{\sum_{j=1}^2 \sqrt{\sum_{l=1}^{h_i} \frac{(v_{1ijl} - v_{2ijl})^2}{(v_{1ijl} + v_{2ijl})^2}}}{2}, \text{ where } i = 2, 3, 4 \quad (8)$$

w_{1_1} , w_{1_2} , w_{1_3} and w_{1_4} are the weights of the node T, DS, FQ, and FT, respectively, related to agent 1, which indicate the importance of these nodes in comparing two Bayesian networks. c_1 , c_2 , c_3 and c_4 are the results of comparing agent 1 and agent 2’s CPTs about node T, DS, FQ and FT. Since the node T is the root node and it has only one column in its CPT, while other nodes (DS, FQ, FT) are the leaf nodes and have two columns of values in their CPTs, we compute c_1 differently from c_2 , c_3 , and c_4 . h_i denotes the number of values in the corresponding node. $h_2 = 3$; $h_3 = 3$; $h_4 = 5$. v_{111} and v_{112} are the values of $p(T = 1)$ and $p(T = 0)$ related to agent 1. v_{211} and v_{212} are the values of $p(T = 1)$ and $p(T = 0)$ related to agent 2. v_{1ijl} and v_{2ijl} are the values in agent 1’s CPTs and agent 2’s CPTs, respectively.

The idea of this metric is that agents compute not only their trust values, their CPTs, but also take into account their preferences (encoded as the weights, w_{1_1} , w_{1_2} , w_{1_3} , w_{1_4}). So agents with similar preferences, such as the importance of file type, quality, download speed, will weight each other’s opinions higher.

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 an agent. Every agent only knows other agents directly connected with it and a few file providers at the beginning. In Figure 2, the circles stand for agents and the rectangles denote file providers.

Every agent 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 agent’s interests in the corresponding file type. The files the agent wants to download are generated based on its interest vector. Every agent keeps two lists. One is the agent list that records all the other agents that the agent has interacted with and its trust values in these agents. The other is the file provider list that records the known file providers and the corresponding Bayesian networks representing the agent’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.




Figure 2. The Network Structure

Our experiments involve 10 different file providers and 40 agents. Each agent will gossip with other agents periodically to exchange their Bayesian networks. The period is 5, which means after each 5 interactions with other agents, the agent will gossip once. $w_{ds} = w_{fq} = 0.5$; $a = 0.3$; $b = 0.5$; $w_{1_1} = w_{1_2} = w_{1_3} = w_{1_4} = 0.25$. The total number of interactions is 1000. We run each configuration for 10 times and use means for the evaluation criteria.

5. Results

The goal of the first experiment is to see if a Bayesian network-based trust model helps agents to select file providers that match better their preferences. Therefore we compare the performance (in terms of percentage of successful recommendations) of a system consisting of agents with Bayesian network-based trust models and a system consisting of agents (without Bayesian networks, BN) that represent general trust, not differentiated to different aspects. Successful recommendations are those positive

recommendations (obtained based on formula 3) when agents are satisfied with interactions with recommended file providers. If an agent gets a negative recommendation for a file provider, it will not interact with the file provider. We have two configurations in this experiment:

- Trust and reputation system with BN: the system consists of agents with Bayesian networks-based trust models that exchange recommendations with each other;
- Trust and reputation system without BN: the system consists of agents that exchange recommendations, but don't 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.

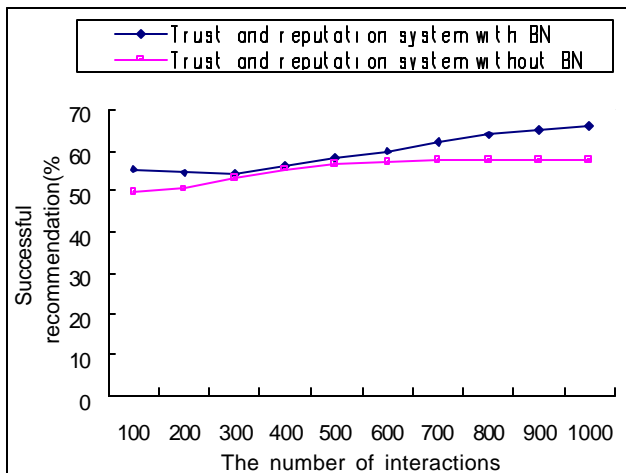


Figure 3. Trust and Reputation System with BN vs. Trust and Reputation System without BN

The goal of the second experiment is to see if exchanging recommendation values with other agents helps agents to achieve better performance (defined as the percentage of successful interactions with file provider). 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 agents with Bayesian networks-based trust models, which don't exchange recommendations with each other;
- Trust system without BN: the system consists of agents that have no differentiated trust models and don't exchange recommendations with each other.

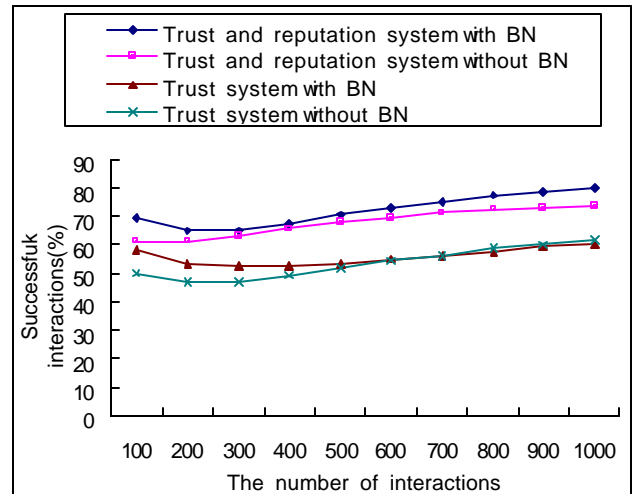


Figure 4. The Comparison of Four Systems

Figure 4 shows that the two systems, where agents share information with each other, outperform the systems, where agents 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 agents 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, an agent's Bayesian network can be viewed as the model of a specified file provider from the agent's personal perspective. In our experiments, we use a very simple naive Bayesian network, which can not 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. If we build a more complex Bayesian network and add more aspects into it, the system performance might be improved.

6. Discussion and Related work

How many Bayesian networks can an agent afford to maintain to represent its trust in other agents in the networks? It depends on the size of the network and the likelihood that agents have repeated interactions. Resnick [15] 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 rare. 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.

Our approach is useful in situations where two agents 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 [9] proves that the small-world phenomenon occurs in Gnutella. It means that agents are inclined to get files from other agents from a small sub-community. This small sub-community often consists of agents that have similar preferences and viewpoints.

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 [3, 5, 11].

Sabater and Sierra's work [16] 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 an agent is compositional. The overall reputation is obtained as a result of the combination of the agent's reputation in each aspect. Our approach integrates these two previous works [1, 16], 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 trusts 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 an agent 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 an agent 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 [13, 20].

Cornelli's work [5] is also in the area of file sharing in peer-to-peer networks. However, it concentrates on how to prevent the attacks to the reputation system and does not discuss how agents model and compute trust and reputation.

7. Conclusions

In this paper, we first review trust and reputation mechanisms from different perspectives. Then 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 agents communicate their experiences (recommendations) outperforms the system where agents 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 an agent can locate a trustworthy service provider. Applying this approach to peer-to-peer systems for computational services is particularly promising.

Acknowledgements

We are grateful to anonymous reviewers for their feedbacks. This research is funded by NSERC Individual Discovery Grant to 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] Abdul-Rahman A. and Hailes S. "A Distributed Trust Model". In Proceedings of the 1997 New Security Paradigms Workshop, 48-60. ACM, 1997.
- [3] Azzedin F. and Maheswaran M. "Evolving and Managing Trust in Grid Computing Systems". IEEE Canadian Conference on Electrical & Computer Engineering (CCECE '02), May 2002.
- [4] Carter J., Bitting E. and Ghorbani A. "Reputation Formalization for An Information-Sharing Multi-Agent System", 515-534. Computational Intelligence, Volume 18, Number 4, November 2002.
- [5] 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.
- [6] Esfandiari B. and Chandrasekharan S. "On how agents make friends: Mechanisms for trust acquisition". In 4th Workshop on Deception, Fraud and Trust in Agent Societies, Montreal, Canada, 2001.
- [7] Falcone R. and Shehory O. "Trust Delegation and Autonomy: Foundations for Virtual Societies". AAMAS tutorial 12, July 16, 2002.
- [8] Hales, D. "Group Reputation Supports Beneficent Norms". The Journal of Artificial Societies and Social Simulation (JASSS) vol. 5, no. 4, 2002.

- [9] Jovanovic M. "Modeling Large-scale Peer-to-Peer Networks and a Case study of Gnutella", University of Cincinnati, master thesis, April 2001.
- [10] Milojicic D. S., Kalogeraki V. and Lukose R. "Peer-to-Peer Computing", Tech Report: HPL-2002-57, <http://www.hpl.hp.com/techreports/2002/HPL-2002-57.pdf>
- [11] Montaner M. and López B. "Opinion based filtering through trust". In Proceedings of the 6th International Workshop on Cooperative Information Agents (CIA'02), Madrid (Spain), September 18-20 2002.
- [12] Mui L., Halberstadt A. and Mohtashemi M. "Notions of reputation in multi-agents systems: A review". In Proceedings of Autonomous Agents & Multiagent Systems (AAMAS'02), 280–287, Bologna, Italy, 2002..
- [13] Niu X., McCalla G., Vassileva J. (to appear) "Purpose-based User Modelling in a Multi-agent Portfolio Management System". In Proceedings of User Modeling UM03, Johnstown, PA, June 22-26, 2003.
- [14] Heckerman, D. "A Tutorial on Learning with Bayesian Networks", Microsoft Research report MSR-TR-95-06, 1995.
- [15] 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.
- [16] Sabater J. and Sierra C. "Regret: a reputation model for gregarious societies". In 4hWorkshop on Deception, Fraud and Trust in Agent Societies, 2001.
- [17] Schafer B.J. Konstan A, J, and Riedl J. "Recommender Systems in ECommerce". ACM Conference on Electronic Commerce (EC-99), November 3-5, 1999, Denver, CO.
- [18] Tang T.Y, Winoto P. and Niu X. "Who can I trust? Investigating trust between users and agents in a multi-agent portfolio management system". AAI-2002 Workshop on Autonomy, Delegation, and Control: From Inter-agent to Groups. Edmonton, Canada.
- [19] Vassileva J., Breban S. and Horsch M. "Agent Reasoning Mechanism for Long-Term Coalitions Based on Decision Making and Trust". Computational Intelligence, Vol. 18, no. 4, 2002.
- [20] Vassileva J., McCalla G. and Greer J. (accepted 17 October 2001) "Multi-Agent Multi-User Modeling", to appear in User Modeling and User-Adapted Interaction.
- [21] Yolum P. and Singh M. "Locating Trustworthy Services" In Proceedings of the First International Workshop on Agents and Peer-to-Peer Computing (AP2PC), 2002.
- [22] Yu B. and Singh P. M. "A social mechanism of reputation management in electronic communities". In Proceedings of Fourth International Workshop on Cooperative Information Agents, 154–165, 2000.
- [23] Yu B. and Singh P. M. "An Evidential Model of Distributed Reputation Management". In Autonomous Agents & Multiagent Systems (AAMAS'02), 294–301, Bologna, Italy, 2002.
- [24] 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.