Building trust in peer-to-peer systems: a review

Bo Zhu and Sushil Jajodia

Center for Secure Information Systems,
George Mason University,
Fairfax, VA 22030-4444, USA
E-mail: bzhu@gmu.edu
E-mail: jajodia@gmu.edu

Mohan S. Kankanhalli*

School of Computing,
National University of Singapore,
Singapore 117543
E-mail: mohan@comp.nus.edu.sg
*Corresponding author

Abstract: The decentralised, cooperative and self-organising nature of Peer-to-Peer (P2P) systems help to mitigate and even overcome many challenges which overwhelm the traditional client-server approaches. On the other hand, these very characteristics also introduce some novel issues in P2P environments. One of the critical issues is how to build the trust relationship within P2P systems. In this paper, we first discuss the desired properties that need to be considered while building trust in P2P systems. Then, we analyse two types of attacks–both the ones mitigated by as well as the ones aimed at trust systems. After this, we divide the previous research work on building trust in P2P systems into two broad categories, that is, reputation-based and trade-based. We then review and discuss the advances in this area based on this classification. Finally, we point out some potential research directions in building trust securely in P2P systems.

Keywords: Peer-to-Peer (P2P) systems; trust; reputation-based trust schemes; trade-based trust schemes.

Reference to this paper should be made as follows: Zhu, B., Jajodia, S. and Kankanhalli, M.S. (2006) 'Building trust in peer-to-peer systems: a review', *Int. J. Security and Networks*, Vol. 1, Nos. 1/2, pp.103–112.

Biographical notes: Bo Zhu is a Post-doctoral Researcher with Center for Secure Information Systems at the George Mason University. He received a PhD in Computer Science and an MSc in Information System from National University of Singapore in 2006 and 2002, respectively. Before that, he received an MEng and a BEng in Automation from Wuhan University in 1999 and 1996, respectively. His research interests include security and privacy in various types of networks (ad hoc/sensor/peer-to-peer/wireless/Internet/grid), digital signatures with various properties, system security and intrusion detection.

Sushil Jajodia is a BDM International Professor of Information Technology and the Director of Center for Secure Information Systems at the George Mason University, Fairfax, Virginia. His research interests include information security, temporal databases and replicated databases. He has authored five books, edited 24 books and published more than 275 technical papers in the refereed journals and conference proceedings. He received the 1996 Kristian Beckman award from IFIP TC 11 for his contributions to the discipline of Information Security. He has served in different capacities for various journals and conferences.

Mohan Kankanhalli obtained a BTech in Electrical Engineering from the Indian Institute of Technology, Kharagpur and an MS and a PhD in Computer and Systems Engineering from the Rensselaer Polytechnic Institute. Science, Bangalore. He is a Professor in the School of Computing at the National University of Singapore. He is on the editorial boards of several journals. His current research interests are in Multimedia Systems (content processing, retrieval) and Multimedia Security (surveillance, watermarking and authentication).

1 Introduction

The surge of *Peer-to-Peer* (*P2P*) networks has swept over the world in the last several years. Many P2P commercial products and free software packages are available for different purposes, for example, file sharing, content distribution, cooperative computing. P2P communication has occupied a large part of internet traffic. From 1999 onwards, Napster (1999) has offered a platform for file sharing which generated more than 20% of traffic on IP networks in the USA within the first few months. According to the results of recent surveys by several internet service providers (Azzouna and Guillemin, 2004; Parker, 2004), more than 50% of internet traffic is due to P2P applications, sometimes even more than 80%. Moreover, it is believed that this trend would continue in the next decade. The rapid increase in both the network bandwidth and computer resources (both CPU power and storage) provides the platform for more powerful and complex functionalities of future P2P systems. As a result, there has been a general trend of extending the inherent P2P characteristics, for example, self-organisation and decentralised architecture, to different environments or upon different objects. For example, the *Mobile Ad hoc Network* (*MANET*) can be viewed as a P2P network in the mobile and wireless environment, and P2P streaming internet TV, for example, CoolStreaming (Zhang et al., 2005), is a P2P technology for sharing live streaming media.

The decentralised, cooperative and self-organising nature of P2P systems help to overcome or at least mitigate many challenges, in which the traditional client–server approaches fail or are inefficient. For example, using the traditional client–server approaches, the server side becomes the bottleneck when the number of clients increase. In contrast, in a P2P system, each user plays the role of both the server and the client which makes the system a lot more scalable. In addition, the central server could be the single point of failure, due to either physical problems or intentional attacks, for example, the *Denial-of-Service* (*DoS*) attacks. The decentralisation of the functionalities of central servers helps mitigate the risks of potential failures of the whole system.

On the other hand, these very characteristics (especially decentralisation) introduce a few new issues in P2P environments. One of the critical issues is how to build trust relationship within P2P systems. Note that, unlike trust in cryptographic systems, where trust usually means the authenticity of an entity, *trust* that we discuss in this paper has broader implications. Informally, for the purpose of this paper, trust means the confidence that a peer has to ensure that it will be treated fairly and securely, when interacting with another peer, for example, during transactions or downloading files. In Section 3.2.7, we provide a more detailed explanation of trust through elaboration of the properties that the P2P trust schemes should possess.

In some P2P scenarios like e-commerce applications, peers are highly dynamic. In particular, when the size of the network increases, the chance that a given pair of peers have repeated interactions with each other is small. Hence, the transaction parties may not have prior experience and knowledge about each other, and peers have to find a way to evaluate the risk involved in the transaction. In other types of P2P applications, for example, file sharing, peers not only need to estimate the trust level of the source providing files to prevent themselves from downloading fake or malicious contents, but they also evaluate the credibility of the peer requesting files to ensure fairness.

Due to the decentralised nature, trust establishment in P2P systems have to necessarily rely on the collaboration among all of the members. Unfortunately, in the unregulated and uncontrolled internet environment, there may exist different kinds of malefic users. Some may try to execute malicious operations to obtain monetary benefits, while some may even launch attacks on the system just for fun. Even worse, in some cases, for example, freeriding, the percentage of malicious users could be very high. Moreover, previous research (Ba and Pavlou, 2002; Houser and Wooders, to appear; Lee et al., 2000) shows that trust-based reputation or recommendation systems have an economic impact, for example, affecting the prices or the sale of products for the e-commerce situation. As a result, unscrupulous producers may try to cheat the system to boost the sales their products or slander the products of their competitors. Furthermore, without central control in P2P systems, there is a lack of accountability, and the fact that many current P2P applications like Freenet (Clarke et al., 2001) and Tangler (Waldman and Mazières, 2001) are designed to provide peer anonymity makes the problem even more challenging.

The purpose of this paper is to review past research work in building trust in P2P systems, to summarise the advances so far, and to point out potential future work. The outline of the paper is as follows. In Section 2, we present the properties that need to be considered in the P2P trust schemes. Then, we enumerate both the known attacks mitigated by and those aiming at the P2P trust schemes in Section 2.5. Afterwards, in Section 3.2.7, we classify trust schemes into two categories, that is, reputation-based and trade-based, and review the existing literature. Finally, we draw the conclusion and point out a few potential future research directions in Section 4.3.

2 Properties of P2P trust schemes

In this section, we discuss about the desired properties, or the design goals, that should be considered when designing trust schemes in the P2P environments.

2.1 Types of feedback

Trust can be built upon the past good and bad experiences, in form of positive and negative feedback, from other peers. Trust schemes that depend on only one type of feedback are inadequate.

The trust systems based on past positive feedbacks only can be cheated in a way that, colluding peers send good reports for each other, for example, claiming other colluding members are storing many files for itself. One way to detect such misbehaviour is to query the virtual/non-existent files

for the purpose of validation. However, if the query is done regularly with certain probability, the overhead would be high. Moreover, this method is suitable for applications for file sharing, but would fail in transaction-based P2P applications. On the other hand, considering only negative feedback is insufficient as well, because newly joined peers would be rated the same or better than peers that consistently perform well (Friedman and Resnick, 2001), which is obviously unfair.

2.2 Authentication and non-repudiation

Basically, the trust system for P2P networks is built upon feedback from numerous peers. We need to make sure that the feedback is really from the source that it claims to be, that is, authentication, so that in case that the feedback is questionable we should be able to find out the peer that is responsible for the fake information. Besides that, the peer detecting and accusing another peer of a fake feedback should be able to provide proof that the source of the fake feedback cannot be denied and at the same time can be verified by other peers, that is, non-repudiation. Authentication and non-repudiation are a must in order to achieve accountability in any trust scheme.

2.3 Communication and storage cost

Naturally, decentralisation of P2P systems leads to more information being exchanged among peers in comparison to the amount exchanged between the server and clients in centralised systems. Besides that, due to the anonymity requirement, certain amount of bogus traffic is always generated to cover the trace of the real traffic.

In trust schemes, the peer querying the trust value of another peer needs to collect enough feedback before determining whether to make transactions. Ideally, the query can be flooded in the P2P network, so as to obtain the most accurate image of the trust level of the queried peer. Unfortunately, without a central party, the overhead of flooding is unacceptable, especially when the size of the network is huge. Intuitively, there is a trade-off between spending more on communication costs and achieving better assurance about the credibility of the potential transaction party. Similarly, there is also a trade-off between storing more feedback to have a higher assurance about the credibility of the potential transaction party and saving of storage space which can be utilised for other purposes.

2.4 Scalability

The scalability is tightly related to the communication and storage costs discussed above. In anonymous P2P systems, to ensure the anonymity of the sender or the receiver, the size of the system is expected to be very large. Current popular P2P applications, for example, Skype (2003), are believed to have millions of users, and hundreds of thousands users might be simultaneously online (Garfinkel, 2005). As a result, the trust scheme is still expected to be efficient while the system scales up.

2.5 Anonymity

Anonymity is an important consideration in designing P2P networks. There have been many protocols proposed to preserve the anonymity of peers (Dingledine, 2004; Reed et al., 1998; Reiter and Rubin, 1998; Shields, 1999). The common theme of these works is to protect the identity of either the sender or the receiver.

Anonymity in terms of trust schemes refers to the protection of the identity of the peer which provides feedback against another peer. One motivation to provide such anonymity is that the peer which gives a feedback does not want the others to know that it has previous interactions with the node being queried. Another motivation is that we wish to hide the trust topology of the network. When using transitive trust, the trust topology is open to all, and malicious nodes could make use of such precious information to choose the victim that could bring the largest benefits. Moreover, maintaining anonymity protects peers sending feedback from being a target of vindictiveness.

3 Attacks related to P2P trust schemes

Previous experience shows that anonymity of the P2P systems gives rise to the possibility of misuse and abuse by malicious peers. For example, they could use the P2P system to distribute malicious (e.g. Trojans and viruses) or illegal (e.g. pornography) content. Instead of introducing all the attacks against P2P systems, in this section we focus on attacks that are either mitigated by or aimed at the trust schemes.

3.1 Attacks mitigated by trust schemes

3.1.1 Freeride

Freeriding (Adar and Huberman, 2000) is a type of uncooperative behaviour in which some users only consume resources of other members in the network without ever contributing as a pay-back. Freeriding nodes, called freeriders, may give incorrect responses to requests from others. They may report falsely about their bandwidth capacity so as not to have much traffic routed over them (Saroiu et al., 2002). Recent studies of Gnutella and Napster confirm that many users consume without contributing (Adar and Huberman, 2000; Saroiu et al., 2002). In particular, a study of the Gnutella file sharing system shows that almost 70% of the peers only consume resources but do not provide any file (Adar and Huberman, 2000).

Without any mechanism to control, the Freeride attack may undermine the cooperative nature of P2P systems, and affect the effectiveness of normal functionalities of the systems. For example, since the anonymity of P2P systems is tightly related to the average number of peers in the system (Guan et al., 2002; Levine and Shields, 2002; Serjantov and Danezis, 2002). Freeride directly downgrades the basic quality of the anonymity service.

It is believed that trust schemes can help lower the extent of freerides. The freeriders will be assigned a lower trust value, which results in less cooperation from others. In other words, the trust scheme works as an incentive mechanism.

3.1.2 Pollution

Pollution is a kind of attack that tries to add bogus files, usually with the same titles but modified contents of the most popular files, into the P2P file sharing system. Pollution attacks can be classified into two major categories (Liang et al., 2005): Content Pollution and Metadata Pollution. Currently, the former is the more common form of pollution. A typical example of this form of attack is to add tens of seconds of undecodable white noise into the middle of a popular song. The latter may tamper the metadata of a file instead of its content. Taking the same example, the polluting party may change the song title or the artist name of a popular song.

Although peers can detect the pollution by matching or user filtering (Liang et al., 2005) after they download the files, in such cases the bandwidths of peers have been wasted anyway. Therefore, some type of a predownloading solution would be useful. One possible solution is to make a prejudgement based on the trust of the peer providing the file.

3.1.3 Worms, viruses and trojans

The open nature of the P2P system makes it a perfect platform for the propagation of malicious programs, that is, worms, viruses and trojans. There are several countermeasures to the threat of P2P worms. One is to use type-safe languages, like Java and C#, to write P2P client programs. Another countermeasure is to increase the diversity of both the client programs and the underlying platforms to limit the propagation rate. Trust schemes, for example, XREP (Damiani et al., 2002), can help mitigate the risks by building resource-based reputation.

3.2 Attacks aimied at trust schemes

3.2.1 Sybil attack

In the sybil attack (Douceur, 2002), malicious peers can present multiple identities, and thus can control a substantial part of the system. If such an attack is possible, it would undermine the basis of the trust schemes, that is, most of peers are honest and each peer can contribute only once, for example, each peer can vote only once, when assessing a given peer.

3.2.2 Denial of service attacks

There are two kinds of denial-of-service attacks against trust schemes. The more common one is that, in a transaction-based system, malicious nodes can flood numerous fake feedback messages through fake transactions (Srivatsa et al., 2005). Another type is that, when malicious nodes detect negative feedback against themselves, as a revenge, they may launch out-of-band (in terms of overlay reputation networks) denial-of-service attacks to disable the normal functionalities of the peers sending the feedback.

3.2.3 False accusation

In P2P systems, it is possible that malicious peers send false accusations, or provide false reports, against an innocent peer.[1] An intuitive method for prevention of false accusation is to collect several reports from different peers before judging the guilty of a given peer.

3.2.4 Context-based attacks

Most existing P2P trust schemes do not support context-based factors while evaluating the trust values of peers. Considering the case that different transactions have different values or weights, malicious peers would choose to be honest in a large number of small-valued transactions, and then try to cheat in large-valued transactions.

3.2.5 Strategic dynamic personality attacks

Most of the existing P2P trust schemes use a combination of average feedback value and the number of transactions performed by a node as indicators of its trust value. For example, in eBay's reputation system, the format of published feedback is the sum of positive, negative and neutral ratings received during a given time period, for example, one week, or one month, or six months (Dellarocas, 2003; Dellarocas et al., 2004). In such a system, malicious peers can build a reputation and then start cheating or oscillating between building and milking the reputation.

3.2.6 Shilling attack

As the trust scheme used becomes a factor that may affect the price or the sale of products, crooked producers may find it profitable to *shill* the system (Lam and Riedl, 2004) in such a way that their items are recommended to users more often than those of their competitors, whether or not their products are of high quality. There have been several real cases of shilling attacks. Amazon.com has detected precisely such a type of attack manipulating its new feature aimed at offering customers a wider range of buying tips. It had to deactivate the 'What's Your Advice?' feature because of unexplained 'commercial abuse' (Dotinga, 2002). Similarly, eBay changed its feedback policy to prohibit users from buying or trading feedback from other members (Steiner, 2003).

3.2.7 Collusion

To be more accurate, collusion itself is not an independent kind of attack. Instead, it can be combined with and enhance the attacks discussed previously. The reason that we list it separately is due to the fact that, although some attacks performed by single node is easy to detect by current reputation mechanisms, collusion between malicious peers make the detecting job much more difficult.

4 Trust in P2P systems

As a result of no central control and monitoring, fairness is a fundamental requirement of a P2P system, especially

for the e-commerce application. The absence of fairness may result in significant degradation of service performance and unpredictable availability of the resources, or even the collapse of the whole P2P system. Unfortunately, previous research by Hardin (1968) in social science indicates that people tend to abuse shared resources that they do not need to pay, and this observation has been proved to be correct in P2P systems (Adar and Huberman, 2000). Consequently, certain level of trust should be built up to give confidence to members in the P2P systems that they will be treated fairly and that they can trade in a fair manner. Accordingly, many trust schemes, or incentive mechanisms and accountability methods, are proposed to build up the trust and ensure the fairness in P2P systems. We classify all of these schemes into two categories:

Reputation-based: reputation-based trust schemes are based on the idea of 'word-of-mouth', that is, a peer's reputation is determined by other peers' opinions. In reputation-based trust schemes, the trust level of a peer denoted as A is calculated using a trust metric which combines different factors, for example, current feedback from other peers towards A, past feedback as well as the context-based information about each feedback.

Trade-based: the trade-based trust schemes are based on the idea of fair exchange, where a peer contributing to other peers is explicitly remunerated, either directly or indirectly. Taking content distribution as an example, when a peer denoted as A helps another peer denoted as B to store one file, it may also require B to reserve the same amount of storage for its usage.

4.1 Reputation-based trust schemes

The main challenges of building reputation-based trust schemes include:

1 how to incorporate necessary information to ensure the accuracy of the reputation
2 how to detect or prevent various attacks from malicious peers and
3 how to ensure the efficiency while the system scales up.

Research on how to evaluate the truth in open networks (Beth et al., 1994; Zimmermann, 1994), like internet, began in the early 1990s. Abdul-Rahman and Hailes (1997) propose a distributed trust model, which includes both trust generalisation and a recommendation protocol. This landmark paper builds up a general framework on which many reputation-based P2P trust schemes could fit in.

Similar to Abdul-Rahman and Hailes (1997), Beth et al. (1994), Zimmermann (1994), Dellarocas (2000), Manchala (2000) have proposed schemes for trust evaluation in centralised e-commerce, but the core ideas along with the issues addressed can be extended into P2P environments. Dellarocas (2000) analyses unfair behaviour in online trading communities, and proposes a set of mechanisms, for example, cluster filtering, to reduce the negative effects of such fraudulent behaviour. Manchala (2000) proposes several trust models, based on boolean relations or fuzzy logic or transaction processes, to perform risk analysis. It also addresses the issue of trust propagation.

One of earliest works for trust evaluation in P2P systems is Aberer and Despotovic (2001). The trust model proposed is based on binary trust, that is, a peer is either trustworthy or not. One weakness of Aberer and Despotovic (2001) is that complaints are the only behavioural data used in the model, and thus a newly joined peer has the same level of reputation as that of a trustworthy peer which has had many transactions performed honestly.

Cornelli et al. (2002) propose P2PRep on top of Gnutella to estimate the trustworthiness of a node. P2PRep consists of two solutions: basic polling and enhanced polling. In the former, the servents[2] responding to the poll do not provide their servent identities denoted as $ServentID$. In the enhanced polling, voters also declare their $ServentID$s so that the weighting of the votes received, that is, the querying node's opinion on the credibility of the voters, can also be taken into account. Although many later trust value-based protocols, for example, (Kamvar et al., 2003; Xiong and Liu, 2003), is based on ideas similar to enhanced polling, in Cornelli et al. (2002) the authors do not discuss trust metrics explicitly.

Unlike (Aberer and Despotovic, 2001; Cornelli et al., 2002) which only consider the reputation of the entity holding the resource, XRep (Damiani et al., 2002) proposed by Damiani et al. is the first work that uses combined reputations of servents and resources, providing more informative pollings and overcoming the limitations of only servent-based solutions. As a result of double pollings, XRep can efficiently detect several attacks, for example, self replication, pseudospoofing and shilling. XRep supports weak anonymity for the peer. In other words, the reputation is bound to a pseudonym or say opaque identifier, that is, the digest associated with a servent, although the real IP address of the peer is still required when it replies to a voting query.

Kamvar et al. (2003) propose the EigenTrust algorithm similar to PageRank (Page et al., 1998) which assigns each peer a unique global trust value, based on the peer's history of uploads. This approach can decrease the number of downloads of inauthentic files in a P2P file-sharing network. In addition, based on the global trust value, the network can identify malicious peers and isolate them from the network. Compared to XRep (Damiani et al., 2002), one major difference is that EigenTrust takes into consideration the trustworthiness of peers which reply to the query for trust evaluation. This change may result in a more accurate trust value. However, the calculation of the global trust value discloses the trust topology of the whole system, which may facilitate the attacks of malicious nodes. In particular, malicious nodes may choose a peer with the highest trust value with all its neighbours as the target for compromise. Another weakness of EigenTrust is that the algorithm relies on a set of well-known pre-trusted users. In addition, for the EigenTrust scheme, to let the trust matrix become a Markov matrix and the global trust rating computation converge to the principal eigenvector of that matrix, the local trust values are normalised, and it results in significant loss of trust information (Selcuk et al., 2004). For example, if there are n identical trust ratings in the database, the normalised value of them will be $1/n$ whether the original values were the highest

possible ratings or the lowest; or a single non-zero rating in the database will always be normalised to 1, regardless of its value.

Guha et al. (2004) develop a framework of trust propagation schemes, and evaluate their schemes on a large, real world, working trust network from the Epinions website. The method of trust propagation in Guha et al. (2004) is similar to EigenTrust (Kamvar et al., 2003) but with a few major differences. Previous experience with real-world implemented trust systems such as Epinions and eBay suggests that distrust is at least as important as trust. Guha et al. (2004) is the first paper that proposes to incorporate distrust in a computational trust propagation setting, while EigenTrust only addresses the positive trust. Another difference is that, the goal of EigenTrust is to assign to each node an universal measure of trust, while the major goal of (Guha et al., 2004) instead is to predict an unknown trust/distrust value between any two users. Moreover, in Guha et al. (2004), the authors define a few types of new trust propagations, that is, cocitation, transpose trust and trust coupling, besides direct propagation used in EigenTrust (Kamvar et al., 2003).

Xiong and Liu (2004) design a reputation-based trust supporting framework–PeerTrust. They introduce three basic trust parameters, that is, feedback a peer receives from other peers, the total number of transactions a peer performs, the credibility of the feedback sources and two adaptive factors, that is, transaction context factor and community context factor, in computing trustworthiness of peers. They also define a general trust metric to combine these parameters. The main contribution of their work is to design a comprehensive framework covering most of important factors that affect the reputation/trust of a peer. In addition, the authors notice that previous trust value-based metrics like EigenTrust (Kamvar et al., 2003), denoted as *TVM* in Xiong and Liu (2004), are based on two assumptions:

1 untrustworthy peers have a higher probability of submitting false or misleading feedback in order to hide their own malicious behaviour
2 trustworthy peers are believed to be honest with a high probability on the feedback they provide.

Xiong and Liu argue that, although the first assumption is true in most cases, the second one is not always true. Therefore, they propose another trust metric based on the querying peer's personalised experience, which is denoted as *PSM* in Xiong and Liu (2004). Concretely, let A and B denote the querying peer and the peer replying to the query, respectively. And A queries B for its opinion (or say trust value) towards another peer denoted as C. Let $P = \{p_1, p_2, \ldots, p_n\}$ denote the group of peers that both A and B have interacted previously. Let $T(X, Y)$ denote the trust value that peer X assigns to peer Y. In the *TVM* metric, the weight of $T(B, C)$ is $T(A, B)$. In the PSM metrics, the weight is the similarity between $\{T(A, p_1), T(A, p_2), \ldots, T(A, p_n)\}$ and $\{T(B, p_1), T(B, p_2), \ldots, T(B, p_n)\}$. Simulation results in Xiong and Liu (2004) show that, compared to the *TVM* metric, the *PSM* metric is more effective in defending against the collusion attack.

As an extension of Xiong and Liu (2004), Srivatsa et al. propose TrustGuard (Srivatsa et al., 2005), a highly dependable reputation-based trust building framework. Unlike (Xiong and Liu, 2004), the major goal of TrustGuard focuses on the vulnerabilities of a reputation system itself. The authors identify three types of threats, that is, strategic oscillations, fake transaction and dishonest feedback, and provide corresponding countermeasures. The paper addresses a few issues that are missed in Cornelli et al. (2002), Damiani et al. (2002), Kamvar et al. (2003), that is, the temporal dimension of the reputation systems (strategic behaviour by malicious nodes) and the problem of fake transactions. TrustGuard is resistant to random shilling attacks, but may be vulnerable to other types of shilling attacks (Srivatsa et al., 2005). When describing the countermeasure against fake transactions, the authors claim that their focus is on 'building a distributed and decentralised solution to handle fake transactions'. However, the paper does not give details about the distribution and decentralisation of the trust authority, and their solution actually still relies on the periodically online Trust Third Party (TTP). Moreover, the proofs that act as signed contracts are exchanged before the actual transaction takes place. A possible attack could be, a malicious node stops the transaction after exchanging the proof, and then provides feedback on the transaction which is listed in the proof but actually does not happen.

There are several reputation-based incentive mechanisms (Blanc et al., 2005; Lai et al., 2003; Ranganathan et al., 2003) based on the Prisoner's Dilemma. In Lai et al. (2003), the evolutionary Prisoner's Dilemma is employed for file sharing. Lai et al. study strategies based on private and shared history and strategies that adapt to the behaviour of strangers. In Ranganathan et al. (2003) and Blanc et al. (2005), the Prisoner's Dilemma is used for file sharing and routing, respectively. Blanc et al. (2005) model the interactions between nodes as a 'random matching game', and describe a simple reputation system that provides incentives for good behaviour. They assume the existence of a trusted authority, who observes the players actions and updates their reputations accordingly. In addition, they simulate malicious peers as peers that always defect, and the effectiveness of this protocol against Byzantine peers is not clear. Their reputation system works best when the network is homogenous, that is, all the peers send requests at the same, steady rate and they all choose destinations uniformly at random.

4.2 Trade-based trust schemes

Trade-based trust schemes can be further divided into two sub-categories: currency-based (also called as micropayment schemes) and resource-based. The major difference between these two sub-categories is in that, currency is used as the intermedium for trading resources among peers in the former, while in the latter there is no such intermedium.

4.2.1 Currency-based trading schemes

One of the earliest currency-based trading schemes is MojoNation (2000). As a reward for peers uploading and distributing files, they are given certain amount of payment, in

a form of digital currency called 'Mojo'. 'Mojo' can be used to purchase resources, for example, files, from other peers. The main limitation of this approach is that all transactions had to be cleared in a centralised system, and users were burdened with the management of their Mojo. The designer of MojoNation, Jim McCoy, also realised the problem of centralisation and shut down the MojoNation network in 2002 (McCoy, 2002).

Ioannidis et al. (2002) propose a credential-based network file storage system, Fileteller, and use a micropayment system to pay for both the initial creation of the file and all of the subsequent accesses. This micropayment protocol is based on a trust management system proposed in Blaze et al. (2001). Fileteller itself does not address the malicious peers issue.

Yu and Singh (2003) present an agent-based P2P system, in which each peer is a software agent and the agents cooperate to search the whole system through referral. A static and a dynamic pricing mechanism are proposed to motivate each agent to behave rationally. In the static pricing protocol, the costs for referrals and answers are fixed for all agents, while in dynamic pricing protocol, the querying peer may place different prices for the referrals and answers based on the qualities of services. This system aims at preventing Freeride, and the static pricing protocol is vulnerable to malicious peers, for example, a peer claims to have answers or referrals, but then does not respond after getting the payment. In the dynamic pricing protocol, the authors propose to make use of the previous history to limit such attacks.

Vishnumurthy et al. (2003) propose a monetary scheme called KARMA for file sharing. In KARMA, each peer holds a single scalar value, called *karma*, which can be used to trace the resource consumption and contribution of this peer. *Karma* is monitored by a set of other nodes, called its 'bank set'. These sets of nodes simulate the responsibility of a trust authority, for example, increasing or decreasing a peer's *karma* based on its contribution or consumption. To mitigate the sybil attack, the system limits the rate at which peers can create new identities through the usage of a cryptographic puzzle. The scheme is flexible and decentralised, but has a high complexity and performance overhead.

4.2.2 Resource-based trading schemes

In FreeHaven (Dingledine, 2004; Dingledine et al., 2000), a receipt is generated when two peers want to exchange data. The receipt contains a hash of the public keys for two transaction parties, information about the data traded, and a timestamp. Note that, the receipt in FreeHaven is not a proof of a transaction but rather an indication of a commitment to keep the data received safe. To prevent malicious peers from discarding the data received, the authors propose a 'buddy system' which associates pairs of shares in a given document. The holder of each share will query its buddy periodically to ensure that its buddy is still alive.

Cox and Noble (2003) propose Samsara, an infrastructure for enforcing fairness in P2P storage systems. The authors claim that Samsara is the only system that enforces storage fairness without requiring trusted third parties, monetary payment, certified identities, or symmetric storage relationships. The idea of Samsara is to manufacture symmetric storage relationships where they do not arise naturally. In Samsara, each contributing node creates a claim that the corresponding consuming node must store, and it will check its storage on the consuming one periodically. Since each node is concerned only with the maintenance of its own data and claims, it is suitable for well-structured environments where members inside have static collaboration relationships with each other, for example, distributed storage. On the other hand, the open ad hoc nature of P2P systems is a challenge for Samsara. For example, in a typical P2P network for file sharing, a node denoted as A may contribute to a group of nodes denoted as X, and at the same time consume resources from another group of nodes Y. The members of X and Y need not necessarily be the same. In such a case, the implementation of Samsara may result in the wastage of a large amount of resources. More specifically, members of X have to maintain the same size of storage that they consume from A, but A may not use them or use only a small portion of the reserved storage. This problem is partially addressed in Cox and Noble (2003) in a way that midstream nodes can forward the claim overheads. However, nodes in Samsara prefer retaining claims rather than forward them whenever possible, because they are still responsible for the claim forwarded. If the downstream node cheats, the midstream node will be penalised. Furthermore, Samsara cannot deal with the problem of popular data that are highly replicated by different peers in the network. Moreover, Samsara is targeted towards greedy nodes, and cannot stop malicious nodes who promise to store data and then immediately discard it.

Ngan et al. (2003) propose to use distributed auditing to enforce fair sharing of P2P resources. In this protocol, each node publishes and digitally signs two logs: the local list of files that the local node is storing on behalf of remote nodes, and the remote list of files that other nodes are storing on behalf of the local node. To prevent collusion, the auditor needs to check not only the usage list of the peer under auditing denoted as A but also all the peers reachable from A's local list, and the check is executed recursively. Obviously, the cost of the recursive audit is prohibitively expensive, and thus the authors propose the conduct of random auditing instead.

Unlike (Cox and Noble, 2003; Ngan et al., 2003), SHARP (Fu et al., 2003) can handle with renewable resources such as CPU and bandwidth. Basically, it is not a mechanism for building trust but an accounting and delegation framework for secure distributed resources management. However, the accounting information provided by SHARP could be a very useful source for evaluating the trust.

Instead of typical two-way exchanges, Anagnostakis and Greenwald propose a new incentive mechanism based on n-way exchanges among rings of peers (Anagnostakis and Greenwald, 2004), and present a search algorithm for locating such rings. According to their empirical results, it is sufficient to search for cycles in chains of up to five predecessors. One weakness of n-way exchanges is that the exchange relations among peers including both the identities of peers and the names of objects requested are open to all, which conflicts with the anonymity requirement of many P2P applications. Another weakness is that, an n-way exchange can be built only when all the objects involved in the chain

are available at the same time. Thus, this scheme is less flexible compared to the cash-based schemes (Ioannidis et al., 2002; Yu and Singh, 2003). Besides, an online trusted peer is required as a mediator in each exchange to prevent man-in-the-middle attacks. Furthermore, the cost of communicating the full request tree may be prohibitive for peers with a large number of incoming requests and peers close to them in the request graph. The authors try to control the cost by checking partial request tree instead, namely, trying two-way or three-way exchange first.

4.3 Summary

Neither reputation-based nor trade-based trust schemes can solve all the problems that P2P systems currently face. Both of them have their pros and cons, and are suitable for different environments.

Currency-based trading schemes use a certain currency as an intermedium, and thus the privacy of transactions or exchanges can be readily ensured. In addition, they are free from the 'double coincidence of wants' constraint, that is, the exchanged resources should be available at the same time, on resource-based trading schemes. However, they have the weakness of higher complexity. The schemes must take into consideration many monetary issues, for example, negotiate prices, inflation and deflation of the currency. Besides, all the currency-based trading schemes have to provide start-up funds to new peers, which could be a potential loophole that malicious peers can make use of.

Compared to reputation-based trust schemes, resource-based trading schemes works well in homogenous environments, where peers have similar consumptions. However, due to the 'double coincidence of wants' constraint, they are inappropriate for highly ad hoc environments like P2P file sharing. In particular, when there exist popular objects shared by a large number of peers, the high imbalance in the contribution may result in a large amount of exchanged resources wasted. As a result, currently most of resource-based trading schemes are mainly used for P2P content distribution. Moreover, in most resource-based trading schemes the information about the exchanged resource is open to all for checking the consistency. Thus, privacy is not available in such systems.

Probably because researchers have realised the similarity between the chaos of a P2P system and the disorder in a human society at its initialisation stage, the trust schemes based on the idea of digital 'word-of-mouth', that is, reputation-based trust schemes, are currently the mainstream of research work on building trust in P2P systems. Indeed, without the limitations on trade-based ones, reputation-based trust schemes are more flexible and can be used in all the possible P2P environments, and even some currency-based trading schemes (Ioannidis et al., 2002; Yu and Singh, 2003) themselves have to rely on the reputation systems to help negotiate the prices of resources or services. The major weakness of reputation systems is that, they only become operational after a node has misbehaved. Therefore, it would be better to provide another method during the bootstrapping stage. For example, in a content distributed system, we can implement a distributed replication system like the *Peer-to-peer Information Preservation and Exchange* (*PIPE*) network (Cooper et al., 2002) to protect the availability of the system before the reputation system becomes effective.

5 Conclusion and future work

There have been various types of P2P systems for different purposes, for example, file sharing and content distribution. Each P2P application or system has its own requirements and assumptions. To satisfy these requirements and assumptions, we may need to choose the most appropriate one or even combine different types of trust schemes. One possible future work is to combine the advantages of both trade-based (at the bootstrapping stage) and reputation-based (after the bootstrapping stage) schemes for a P2P file sharing system.

A major reason for the popularity of P2P systems is that they can provide anonymity for the peers. Currently, most of trust schemes concentrate a lot on the accuracy of trust evaluation and the robustness against malicious attacks, but ignore the requirement of anonymity. XRep (Damiani et al., 2002) can provide weak anonymity for the voting peer, but the real IP address of the peer has to be disclosed in order to prevent the pseudo-spoofing attack. How to provide more anonymity in the trust scheme is still an open question.

To prevent from malicious operations and collusion, most resource-based trading schemes require periodical checking on the resource information, for example, the list of exchanged files or resources, which might be relatively costly. Thus, one challenge is to minimise the communication costs involved in periodical checking while the system scales up.

Although there have been extensive research work about protecting trust schemes from various types of attacks, how to accurately and efficiently detect different attacks, for example, fake transactions, in terms of collusion remains an open problem.

Several reputation systems (Guha et al., 2004; Kamvar et al., 2003) assume transitive trust, either direct propagation (Kamvar et al., 2003) or else (e.g. cocitation, transpose trust, trust coupling (Guha et al., 2004)). However, different users have different assessments on the same type of operations, and thus might give different trust values Or they may have different security policy, which would also affect the trust values assigned. Consequently, it is quite challenging to adapt transitive trust under environments with inconsistent assessment settings.

References

Aberer, K. and Despotovic, Z. (2001) 'Managing trust in a peer-2-peer information system', *Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM'01)*, pp.310–317.

Adar, E. and Huberman, B.A. (2000) 'Free riding on gnutella', *First Monday*, September, Vol. 5, No. 10.

Abdul-Rahman, A. and Hailes, S. (1997) 'A distributed trust model', *Proceedings of New Security Paradigms Workshop*.

Anagnostakis, K.G. and Greenwald, M.B. (2004) 'Exchange-based incentive mechanisms for peer-to-peer file sharing', *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS 2004)*, pp.524–533.

Azzouna, N.B. and Guillemin, F. (2004) 'Experimental analysis of the impact of peer-to-peer applications on traffic in commercial IP networks', *European Transactions on Telecommunications: Special Issue on P2P Networking and P2P Services*, Vol. 15, No. 6, pp.511–522.

Ba, S. and Pavlou, P.A. (2002) 'Evidence of the effect of trust building technology in electronic markets: price premiums and buyer behavior', *MIS Quarterly*, Vol. 26, No. 3.

Beth, T., Borcherding, M. and Klein, B. (1994) 'Valuation of trust in open networks', *Proceedings of the Third European Symposium on Research in Computer Security (ESORICS'94), LNCS 875*, pp.3–18.

Blanc, A., Liu, Y. and Vahdat, A. (2005) 'VDesigning incentives for peer-to-peer routing', *Proceedings of IEEE INFOCOM 2005*, pp.374–385.

Blaze, M., Ioannidis, J. and Keromytis, A.D. (2001) 'Offline micropayments without trusted hardware', *Proceedings of Financial Cryptography (FC'01), LNCS 2339*, pp.21–40.

Clarke, I., Sandberg, O., Wiley, B. and Hong, T.W. (2001) 'Freenet: a distributed anonymous information storage and retrieval system', *Proceedings of International Workshop on Design Issues in Anonymity and Unobservability, LNCS 2009*, pp.46–66.

Cooper, B.F., Bawa, M., Daswani, N. and Garcia-Molina, H. (2002) 'Protecting the PIPE from malicious peers', *Stanford InfoLab Technical Report 2002-27*.

Cornelli, F., Damiani, E., De Capitani di Vimercati, S., Paraboschi, S. and Samarati, P. (2002) 'Choosing reputable servents in a P2P network', *Proceedings of the 11th International Conference on World Wide Web*, pp.376–386.

Cox, L.P. and Noble, B.D. (2003) 'Samsara: honor among thieves in peer-to-peer storage', *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, pp.120–132.

Damiani, E., De Capitani di Vimercati, S., Paraboschi, S., Samarati, P. and Violante, F. (2002) 'A reputation-based approach for choosing reliable resources in peer-to-peer networks', *Proceedings of the Ninth ACM Conference on Computer and Communications Security*, pp.207–216.

Dellarocas, C. (2000) 'Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior', *Proceedings of the Second ACM Conference on Electronic Commerce*, pp.150–157.

Dellarocas, C. (2003) 'The digitization of word-of-mouth: promise and challenges of online reputation mechanisms', *Management Science*, October.

Dellarocas, C., Fan, M. and Wood, C.A. (2004) 'Self-interest, reciprocity, and participation in online reputation systems', *MIT Sloan Working Papers No. 4500-04*.

Dingledine, R., Freedman, M.J. and Molnar, D. (2000) 'The free haven project: distributed anonymous storage servic', *Proceedings of Workshop on Design Issues in Anonymity and Unobservability*, pp.59–82.

Dingledine, R., Mathewson, N. and Syverson, P. (2004) 'Tor: the second-generation onion router', *Proceedings of the 13th USENIX Security Symposium*.

Dingledine, R. (2004) 'The free haven project: design and deployment of an anonymous secure data haven', *Master Thesis*, MIT.

Dotinga, R. (2002) 'Amazon pulls plug on 'Advice'', Available at: http://www.wired.com/news/ebiz/0,1272,53634,00.html.

Douceur, J.R. (2002) 'The sybil attack', *Proceedings of the First International Workshop on Peer-to-Peer Systems (IPTPS 2002)*, pp.251–260.

Friedman, E.J. and Resnick, P. (2001) 'The social cost of cheap pseudonyms', *Journal of Economics and Management Strategy*, Vol. 10, No. 2, pp.173–199.

Fu, Y., Chase, J., Chun, B., Schwab, S. and Vahdat, A. (2003) 'SHARP: an architecture for secure resource peering', *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, pp.133–148.

Garfinkel, S. (2005) 'Can 9 million skype users be wrong?' *Computerworld*, Vol. 11, No. 11.

Guan, Y., Fu, X., Bettati, R. and Zhao, W. (2002) 'An optimal strategy for anonymous communication protocols', *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS 2002)*, pp.257–266.

Guha, R., Kumar, R., Raghavan, P. and Tomkins, A. (2004) 'Propagation of trust and distrust', *Proceedings of the 13th International Conference on World Wide Web*, pp.403–412.

Hardin, G. (1968) 'The tragedy of the commons', *Science*, Vol. 162, pp.1243–1248.

Houser, D. and Wooders, J. (to appear) 'Reputation in auctions: theory, and evidence from eBay', *Journal of Economics and Management Strategy*.

Ioannidis, J., Ioannidis, S., Keromytis, A.D. and Prevelakis, V. (2002) 'Fileteller: paying and getting paid for file storage', *Proceedings of Financial Cryptography (FC'02)*.

Kamvar, S.D., Schlosser, M.T. and Garcia-Molina, H. (2003) 'The eigentrust algorithm for reputation management in P2P networks', *Proceedings of the 12th International Conference on World Wide Web (WWW 2003)*, pp.640–651.

Lai, K., Feldman, M., Stoica, I. and Chuang, J. (2003) 'Incentives for cooperation in peer-to-peer networks', *Proceedings of the First Workshop on Economics of Peer-to-Peer Systems (P2PEcon '03)*.

Lam, S.K. and Riedl, J. (2004) 'Shilling recommender systems for fun and profit', *Proceedings of the 13th International Conference on World Wide Web*, pp.393–402.

Lee, Z., Im, I. and Lee, S.J. (2000) 'The effect of negative buyer feedback on prices in internet auction markets', *Proceedings of the 21st International Conference on Information Systems (ICIS '00)*, pp.286–287.

Levine, B.N. and Shields, C. (2002) 'Hordes: a protocol for anonymous communication over the internet', *ACM Journal of Computer Security*, Vol. 10, No. 3, pp.213–240.

Liang, J., Kumar, R., Xi, Y. and Ross, K.W. (2005) 'Pollution in P2P file sharing systems', *Proceedings of INFOCOM 2005*, Vol. 2, pp.1174–1185.

Manchala, D.W. (2000) 'E-commerce trust metrics and models', *IEEE Internet Computing*, Vol. 4, No. 2, pp.36–44.

McCoy, J. (2002) 'MojoNation public network shutting down', Available at: http://marc.theaimsgroup.com/?l=mojonation-users&m=101286320902173&w=2.

MojoNation (2000) Available at: http://www.mojonation.net/Mojonation.html.

Napster (1999) Available at: http://www.napster.com.

Ngan, T.J., Wallach, D.S. and Druschel, P. (2003) 'Enforcing fair sharing of peer-to-peer resources', *Proceedings of the Second International Workshop (IPTPS 2003), LNCS 2735*, pp.149–159.

Page, L., Brin, S., Motwani, R. and Winograd, T. (1998) 'The pageRank citation ranking: bringing order to the web', *Technical Report 1999-66*, Stanford University.

Parker, A. (2004) 'The true picture of peer-to-peer file sharing', Available at: http://www.cachelogic.com/research/slide1.php.

Ranganathan, K., Ripeanu, M., Sarin, A. and Foster, I. (2003) 'To share or not to share: an analysis of incentives to contribute in collaborative file sharing environments', *Proceedings of the First Workshop on Economics of Peer-to-Peer Systems (P2PEcon'03)*.

Reed, M.G., Syverson, P.F. and Goldschlag, D.M. (1998) 'Anonymous connections and onion routing', *IEEE Journal on Selected Areas in Communications, Special Issue on Copyright and Privacy Protection*, Vol. 16, No. 4, pp.482–494.

Reiter, M.K. and Rubin, A.D. (1998) 'Crowds: anonymity for web transactions', *ACM Transactions on Information and System Security (TISSEC)*, Vol. 1, No. 1, pp.66–92.

Saroiu, S., Gummadi, P.K. and Gribble, S.G. (2002) 'A measurement study of peer-to-peer file sharing systems', *Proceedings of Multimedia Computing and Networking 2002 (MMCN'02)*.

Selçuk, A.A., Uzun, E. and Pariente, M.R. (2004) 'A reputation-based trust management system for P2P networks', *Proceedings of IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2004)*, pp.251–258.

Serjantov, A. and Danezis, G. (2002) 'Towards an information theoretic metric for anonymity', *Proceedings of Privacy Enhancing TechnologiesWorkshop (PET 2002), LNCS 2482*, pp.41–53.

Shields, C. (1999) 'Secure hierarchical multicast routing and multicast internet anonymity', PhD Thesis, Computer Engineering, University of California, Santa Cruz.

Skype (2003) Available at: http://www.skype.com.

Srivatsa, M., Xiong, L. and Liu, L. (2005) 'TrustGuard: countering vulnerabilities in reputation management for decentralized overlay networks', *Proceedings of the 14th International Conference on World Wide Web*, pp.422–431.

Steiner, I. (2003) 'eBay tightens feedback loopholes', Available at: http://www.auctionbytes.com/cab/abn/y03/m09/i17/s01.

Vishnumurthy, V., Chandrakumar, S. and Sirer, E.G. (2003) 'KARMA: a secure economic framework for P2P resource sharing', *Proceedings of the First Workshop on Economics of Peer-to-Peer Systems (P2PEcon'03)*.

Waldman, M. and Mazières, D. (2001) 'Tangler: a censorship-resistant publishing system based on document entanglements', *Proceedings of the Eighth ACM Conference on Computer and Communications Security*, pp.126–135.

Xiong, L. and Liu, L. (2003) 'A reputation-based trust model for peer-to-peer e-commerce communities', *Proceedings of IEEE International Conference on E-Commerce (CEC 2003)*, pp.275–284.

Xiong, L. and Liu, L. (2004) 'PeerTrust: supporting reputation-based trust for peer-to-peer electronic communities', *IEEE Transactions on Knowledge and Data Engineering*, Vol. 16, No. 7, pp.843–857.

Yu, B. and Singh, M.P. (2003) 'Incentive mechanisms for peer-to-peer systems', *Proceedings of the Second International Workshop on Agents and Peer-to-Peer Computing (AP2PC'03), LNCS 2872*, pp.77–88.

Zhang, X., Liu, J., Li, B. and Yum, T.P. (2005) 'Cool-streaming/DONet: a data-driven overlay network for efficient live media streaming', *IEEE INFOCOM 2005*, Vol. 3, pp.2102–2111.

Zimmermann, P. (1994) *PGP User's Guide*, MIT.

Notes

[1]Note that, this attack is different from the denial-of-service attack in that the false accusation may be generated towards true transactions.

[2]Nodes in peer-to-peer systems play the role of both server and client. Therefore, the term *servent* is introduced to identify this double responsibility.