

Chapter II

Reputation and Trust

Li Xiong, Georgia Institute of Technology, USA

Ling Liu, Georgia Institute of Technology, USA

Abstract

This chapter introduces reputation systems as a means of facilitating trust and minimizing risks in m-commerce and e-commerce in general. It first illustrates the importance of reputation systems in m-commerce by analyzing a list of risks through example scenarios and discusses a number of challenges of building an effective and robust reputation system in e-commerce applications. It then describes PeerTrust, an adaptive and dynamic reputation based trust model that helps participants or peers to evaluate the trustworthiness of each other based on the community feedback about participants' past behavior. It also presents some initial experiments showing the effectiveness, benefit and vulnerabilities of the reputation systems. Finally it discusses a few interesting open issues.

Introduction

Mobile commerce (m-commerce) communities create enormous opportunities for many, as participants (or peers) can purchase products, access information, and interact with each other from anywhere at any time. However, they also present risks for participants as they are often established dynamically with unknown or unrelated participants. The

open nature of such presents a big challenge for accountability. As in general e-commerce, the participants have to manage the risk when interacting with other participants. In other words, in addition to its wireless communication layer risks, m-commerce is also faced with all the application layer risks in general e-commerce. For example, a Palm Pilot user may encounter a virus attack by downloading the Liberty Trojan masquerading as an innocent program for PalmOS from other malicious users, which will wipe out all the contact information. Techniques such as smart cards solve part of the problem by authentication but cannot answer the question of which players are more trustworthy. It is very important for users to be able to quantify and compare the trustworthiness of different participants so they can choose reliable and reputable ones to interact with and filter out the unreliable ones to reduce risk.

Application and Risk Scenarios

We first analyze the risks through several m-commerce example scenarios and illustrate the importance of reputation based trust systems.

M-commerce communities can be built on top of either traditional client-server architecture or peer-to-peer wireless network. In the first case, mobile devices are connected to fixed networks through a wireless gateway in order to access the services in the wired Internet. It essentially replaces desktop computers with mobile devices in the traditional e-commerce communities and allows users to order products and access information from anywhere and at any time. Several important classes of applications have been identified, including transaction-based applications such as mobile auction and mobile shopping, communication-based applications such as mobile advertising and mobile alerts, and entertainment-based applications such as mobile music and software downloading (Varshney, 2002). M-commerce communities can be also built on top of a P2P network. They are typically formed by a group of mobile devices under the same service coverage that have a common mission or interest. All members or peers communicate over wireless channels directly without any fixed networking infrastructure. Such type of infrastructure is receiving growing attention for commercial applications, such as team collaboration applications, networking intelligent sensors and cooperative robots.

Most m-commerce security techniques or analyses deal with security concerns specific to the wireless communication such as privacy and authenticity of wireless communications (Chari, 2001). However, the application layer risks in general e-commerce are also manifested in m-commerce. Mobile clients or peers have to face potential threats or risks when interacting with unknown or unfamiliar service providers or other peers. We summarize the risks and threats as follows:

- *Transaction Specific Risks.* For example, in mobile auctions scenario, buyers are vulnerable to potential risks because malicious sellers may provide incomplete or distorted information or fail to deliver goods.
- *Malicious SMS Messages.* Applications such as mobile advertising and mobile alerts typically send advertising and alerts to mobile users using short messaging

service (SMS) messages or short paging messages. A malicious service provider or participant may send out malicious SMS messages that hide nefarious instructions.

- *Virus Attack.* Consider the mobile software-downloading scenario where a mobile user is asking for a resource from the network. An adversary can respond by a fake resource with the same name as the real resource the original user is looking for, but the actual file could be a virus. The first wireless virus has been discovered in PalmOS, which is called PalmOS/Phage1. It will infect all third-party applications on the PDA device. Other wireless virus examples include the PalmOS/LibertyCrack2, a Trojan that arrives masquerading as a crack program for an application called Liberty, which allows PalmOS devices to run Nintendo GameBoy Games. When run, however, the Trojan attempts to delete all applications from the handheld and then reboot it.
- *DoS Attack.* The first cell phone virus hacked users of GSM mobile phones and broadcasted a disparaging remark through SMS. Although the virus caused no damage, it foreshadowed a potential DoS attack. If an adversary can disseminate a worm that send out millions of such messages, it could deluge cell phones with them, thereby overwhelming the short message system.

Reputation Systems

Reputation systems (Resnick, 2000) provide attractive techniques to address the above listed risks by facilitating trust and minimizing risks through reputations. Concretely, they help participants to evaluate trustworthiness of each other and predict future behaviors of participants based on the community feedback about the participants' past behavior. By harnessing the community knowledge in the form of feedback, these systems help people decide who to trust, encourage trustworthy behavior, and deter dishonest participation. Reputation systems are important for fostering trust and minimize risks in two ways. First, by collecting and aggregating feedback about participants' past behavior, they provide a way for participants to share their experiences and knowledge so they can estimate the trustworthiness of other participants with whom they may not have personal experiences and in turn they can avoid malicious participants to reduce risk. Second, the presence of a reputation system creates the expectation of reciprocity or retaliation in future behavior, which in turn creates an incentive for good behavior and discourages malicious behavior.

Building such reputation-based systems for m-commerce communities presents a number of challenges. The main one is how to develop an effective trust model that computes an accurate trust value for each participant even with possible strategic malicious behaviors of participants. This essentially applies to general e-commerce communities at large. Dellarocas (2003) provides a latest survey for research in game theory and economics on the topic of reputation. Most of the game theoretic models assume that stage game outcomes are publicly observed. Online feedback mechanisms, in contrast, rely on private (pair-wise) and subjective ratings of stage game outcomes. This intro-

duces two important considerations. One is incentive for providing feedback and the other is the credibility or the truthfulness of the feedback.



A variety of online community sites have reputation management built in, such as eBay, Amazon, Yahoo! Auction, Edeal, Slashdot, and Entrepreneur.com. Even though they facilitate the trust among users to some extent, they also have some common problems and vulnerabilities. Most of these systems use a simple sum or average of the ratings as the reputation value of a user. For example, eBay uses a summation of positive and negative feedback. It fails to convey important subtleties of online interactions such as whether these feedback ratings come from low-value transactions and whether the feedback ratings are honest. It is important to develop effective metrics that aggregate feedback into a meaningful trust value as an estimate of the trustworthiness of participants by incorporating all the subtleties of online interactions. We discuss below the research challenges of developing an effective trust model in detail.

- *Differentiating dishonest feedback.* An important difficulty in aggregating feedback into a single value is dealing with dishonest feedback and various attacks to the reputation system itself. Malicious participants may provide false or misleading feedback to badmouth other participants and to fool the system. Things are made much worse if a group of malicious participants collude to boost each other's ratings and damage others' ratings. An effective trust metric has to differentiate dishonest feedback from honest ones and be robust against various malicious manipulations of participants.
- *Context and location awareness.* Another important consideration is the context and location awareness, as many of the applications are sensitive to the context or the location of the transactions. For example, the functionality of the transaction is an important context to be incorporated into the trust metric. Amazon.com may be trustworthy on selling books but not on providing medical devices.
- *Incentive to provide feedback.* Lastly, there is a lack of incentive for participants to provide feedback. It is even more so in m-commerce communities where mobile users may not bother to provide feedback at all due to the power limitations of their mobile devices and their on-the-road situation.

The other important challenge is related to how to build the supporting infrastructure to collect, aggregate and distribute feedback and reputation information.

- *Efficient and scalable reputation data dissemination.* There are two alternative ways for reputation data dissemination, namely centralized and decentralized. A trust model can be implemented by either scheme. For example, in the m-commerce communities that are built on top of client-server architecture, a centralized trust server (wireless access provider or other independent service provider) can be deployed to collect, aggregate and distribute reputation information. In the peer-to-peer wireless network, the P2P nature of this type of network makes the

traditional centralized solution unfeasible, as there is no centralized server or database. Various P2P data location schemes such as broadcast based scheme and distributed hash table based schemes can be used to store and look up the reputation data. Data replication has to be considered in order to address the dynamics of the network such as members leaving and joining the network and potential malicious behaviors of the peers.

- *Secure trust data transmission.* There are a number of known security threats at the wireless communication layer. The reputation system infrastructure has to guarantee the secrecy and integrity of the reputation data during their transmission. Encryption based wireless security solutions such as WAP WTLS4  PKI5  schemes can be used in the implementation to ensure reputation data are securely transferred.

Bearing these research issues in mind, we developed PeerTrust (Xiong, 2003) as a dynamic and adaptive reputation based trust system for participants or peers to quantify and compare the trustworthiness of each other. The rest of the chapter focuses on the trust model. The next section describes the PeerTrust model. Technical details including the illustration of the trust metrics in the context of e-commerce and m-commerce applications will be provided. The section followed presents some initial experiments evaluating the trust model. The last section concludes the chapter by a summary and points out some future research opportunities.

The Trust Model

The main focus of PeerTrust approach is the design and development of a dynamic trust model for aggregating feedback into a trust value to quantify and assess the trustworthiness of participants or peers in e-commerce communities.

Trust Parameters

A peer's trustworthiness is defined by an evaluation of the peer it receives in providing service to other peers in the past. Such reputation reflects the degree of trust that other peers in the community have on the given peer based on their past experiences. We identify five important factors for such evaluation: (1) the feedback a peer obtains from other peers, (2) the feedback scope, such as the total number of transactions that a peer has with other peers, (3) the credibility factor of the feedback source, (4) the transaction context factor for discriminating mission-critical transactions from less or non-critical ones, and (5) the community context factor for addressing community-related characteristics and vulnerabilities. We now illustrate the importance of these parameters through a number of example scenarios.

Feedback in Terms of Amount of Satisfaction

Reputation-based systems rely on feedback to evaluate a peer. Feedback in terms of amount of satisfaction a peer receives during a transaction reflects how well this peer has fulfilled its own part of the service agreement. Some existing reputation based systems use this factor alone and compute a peer u 's trust value by a summation of all the feedback u receives through its transactions with other peers in the community. For example, buyers and sellers in eBay can rate each other after each transaction (+1, 0, -1) and the overall reputation is the sum of these ratings over the last 6 months.

We can clearly see that these feedback-only metrics are flawed. A peer who has performed dozens of transactions and cheated 1 out of every 4 cases will have a steadily rising reputation in a given time duration whereas a peer who has only performed 10 transactions during the given time duration but has been completely honest will be treated as less reputable if the reputation measures are computed by a simple sum of the feedback they receive. It is been proved that binary reputation mechanisms will not function well and the resulting market outcome will be unfair if judgment is inferred from knowledge of the sum of positive and negative ratings alone (Dellarocas, 2001).

Number of Transactions

As described above, a peer may increase its trust value by increasing its transaction volume to hide the fact that it frequently misbehaves at a certain rate when a simple summation of feedback is used to model the trustworthiness of peers. The number of transactions is an important scope factor for comparing the feedback in terms of degree of satisfaction among different peers. An updated metric can be defined as the ratio of the total amount of satisfaction peer u receives over the total number of transactions peer u has, that is, the average amount of satisfaction peer u receives for each transaction. However, this is still not sufficient to measure a peer's trustworthiness. When considering reputation information we often account for the source of information and context.

Credibility of Feedback

The feedback peer u receives from another peer v during a transaction is simply a statement from v regarding how satisfied v feels about the quality of the information or service provided by u . A peer may make false statements about another peer's service due to jealousy or other types of malicious motives. Consequently a trustworthy peer may end up getting a large number of false statements and may be evaluated incorrectly even though it provides satisfactory service in every transaction.

We introduce the credibility of feedback as a basic trust building parameter, which is equally important as the number of transactions and the feedback. The feedback from those peers with higher credibility should be weighted more than those with lower credibility. We have developed two mechanisms for measuring the credibility of a peer in providing feedback. The concrete formulas will be discussed later.

Transaction Context Factor

Transaction context is another important factor when aggregating the feedback from each transaction as we have discussed earlier because of the context and location awareness of mobile transactions. For example, when a mobile user is trying to compare potential services, the previous feedback from a mobile user who was using the same device and was in the same location to access the service should be weighted more than those from a regular user accessing the service from a desktop computer at home.

Other general transaction context such as the value and functionality are also important. For example, the size of a transaction should be incorporated to give more weight to the feedback from larger transactions. It can act as a defense against some of the subtle malicious attacks, such as when a seller develops a good reputation by being honest for small transactions and tries to make a profit by being dishonest for large transactions. It can be seen as a simplified mechanism for more sophisticated risk management in e-commerce (Manchala, 2000).

Community Context Factor

Community contexts can be used to address non-transaction specific issues. One example is to add a reward for peers who submit feedback. This can to some extent alleviate the feedback incentive problem. As another example, it can be also used to incorporate historical information, and reputation from other applications or communities.

General Trust Metric

We have discussed the importance of each of the five trust parameters. In this section we formalize these parameters, present a general trust metric that combines these parameters in a coherent scheme, and describe the formula we use to compute the values for each of the parameters given a peer and the community it belongs to.

Given a recent time window, let $I(u, v)$ denote the total number of Interactions performed by peer u with v , $I(u)$ denote the total number of interactions performed by peer u with all other peers, $p(u, i)$ denote the other participating peer in peer u 's i th transaction, $S(u, i)$ denote the normalized amount of Satisfaction peer u receives from $p(u, i)$ in its i th transaction, $Cr(v)$ denote the Credibility of the feedback submitted by v , $TF(u, i)$ denote the adaptive Transaction context Factor for peer u 's i th transaction, and $CF(u)$ denote the adaptive Community context Factor for peer u . α and β denote the normalized weight factors, and Trust value of peer u , denoted by $T(u)$, is defined as follows:

$$T(u) = \alpha * \frac{\sum_{i=1}^{I(u)} S(u, i) * (Cr(p(u, i)) * TF(u, i))}{\sum_{i=1}^{I(u)} Cr(p(u, i)) * TF(u, i)} + \beta * CF(u)$$

The first term is a weighted average of amount of satisfaction a peer receives for each transaction. The weight ($Cr(p(u,i))*TF(u,i)$) takes into account the credibility of feedback source to counter dishonest feedback, and transaction context to capture the transaction-dependent characteristics. This history-based evaluation can be seen as a prediction for peer u 's likelihood of a successful transaction in the future. A confidence value can be computed and associated with the trust metric that may reflect the number of transactions, and the standard deviation of the ratings depending on different communities and requirements. The second term adjusts the first term by an increase or decrease of the trust value based on community-specific characteristics. The α and β parameters can be used to assign different weights to the feedback-based evaluation and community context in different situations. For instance, they can be assigned properly so the trust value is set to be either the feedback-based evaluation when the peer has enough transactions or a default value otherwise.

Important to note is that this general trust metric may have different appearances depending on which of the parameters are turned on and how the parameters and weight factors are set. The design choices depend on characteristics of online communities. It is a non-trivial problem to choose the optimal parameters in practice. Different users may also choose different settings based on their own preferences and have their own view of the universe. We emphasize that the first three parameters — the feedback, the number of transactions, and the credibility of feedback source are important basic trust parameters that should be considered in computing a peer's trustworthiness in any communities. We illustrate next how the basic parameters can be collected or determined and how the adaptive parameters can be set.

The Basic Metric

We first consider the basic form of the general metric as shown below by turning off the transaction context factor ($TF(u,i)=1$) and the community context factor ($\alpha=1$ and $\beta=0$). It computes the trust value of a peer u by a weighted average of the amount of satisfaction peer u receives for each transaction.

$$T(u) = \frac{\sum_{i=1}^{I(u)} S(u,i) * Cr(p(u,i))}{\sum_{i=1}^{I(u)} Cr(p(u,i))}$$

The feedback in terms of amount of satisfaction is collected by a feedback system. PeerTrust uses a transaction-based feedback system, where the feedback is bound to each transaction. The system solicits feedback after each transaction and the two participating peers give feedback about each other based on the transaction. Feedback systems differ with each other in their feedback format. They can use a positive format,

a negative format, a numeric rating or a mixed format. $S(u, i)$ is a normalized amount of satisfaction between 0 and 1 that can be computed based on the feedback.

Both the feedback and the number of transactions are quantitative measures and can be collected automatically. Different from these two, the third parameter — credibility of feedback — is a qualitative measure and needs to be computed based on past behavior of peers who file feedback. Different approaches can be used to determine the credibility factor and compute the credible amount of satisfaction. One way is to solicit separate feedback for feedback themselves. This makes the reputation system more complex and adds more burdens on users. A simpler approach is to infer or compute the credibility value of a peer implicitly. We discuss two such credibility measures.

The first one is to use a function of the trust value of a peer as its credibility factor recursively so feedback from trustworthy peers are considered more credible and thus weighted more than those from untrustworthy peers. We refer to the basic trust metric that uses the Trust Value of a peer recursively as its credibility Measure as PeerTrust TVM metric and it is defined as follows:

$$T(u) = \frac{\sum_{i=1}^{I(u)} S(u, i) * T(p(u, i))}{\sum_{i=1}^{I(u)} T(p(u, i))}$$

This solution is based on two assumptions. First, untrustworthy peers are more likely to submit false or misleading feedback in order to hide their own malicious behavior. Second, trustworthy peers are more likely to be honest on the feedback they provide. It is widely recognized that the first assumption is generally true but the second assumption may not be true at all time. For example, it is possible that a peer may maintain a good reputation by performing high quality services but send malicious feedback to its competitors. In this extreme case, using a function of trust value to approximate the credibility of feedback will generate errors. This is because the reputation-based trust in PeerTrust model is established in terms of the quality of service provided by peers, rather than the quality of the feedback filed by peers.

The second credibility measure is for a peer w to use a personalized similarity measure to rate the credibility of another peer v through w 's personalized experience. Concretely, peer w will use a personalized similarity between itself and another peer v to weight the feedback by v on any other peers. Let $IS(v)$ denote the Set of peers that have Interacted with peer v . To measure the feedback credibility of peer v , peer w computes the feedback similarity between w and v over $IS(v) \cap IS(w)$, the common set of peers they have interacted with in the past. If we model the feedback by v and the feedback by w over the common set of peers as two vectors, the credibility can be defined as the similarity between the two feedback vectors. Particularly, we use the root-mean-square or standard deviation (dissimilarity) of the two feedback vectors to compute the feedback similarity. We refer to the basic metric that uses the Personalized Similarity as the credibility Measure as PeerTrust PSM metric and it is defined as follows:

$$T(u) = \frac{\sum_{i=1}^{I(u)} S(u, i) * Sim(p(u, i), w)}{\sum_{i=1}^{I(u)} Sim(p(u, i), w)}$$

where

$$Sim(v, w) = 1 - \sqrt{\frac{\sum_{x \in IS(v) \cap IS(w)} \left(\frac{\sum_{i=1}^{I(x, v)} S(x, i)}{I(x, v)} - \frac{\sum_{i=1}^{I(x, w)} S(x, i)}{I(x, w)} \right)^2}{|IS(v) \cap IS(w)|}}$$

This notion of local or personalized credibility measure provides great deal of flexibility and stronger predictive value as the feedback from similar raters are given more weight. It may also act as an effective defense against potential malicious collusions. Given the observation that peers in a collusive group give good ratings within the group and bad ratings outside the group, the feedback similarity between a peer v in the collusive group and a peer w outside the group will be low, which will effectively filter out the dishonest feedback by peer v for peer w .

Given that one of the design goals of PeerTrust model is to emphasize the roles of different trust parameters in computing trustworthiness of peers, in the rest of the chapter we will use the above two measures as examples and study their effectiveness, benefit and vulnerabilities. We believe that the study of what determines the precision of credibility of feedback is by itself an interesting and hard research problem that deserves attention of its own.

Adapting the Trust Metric with Context Factors

We have discussed the motivations and scenarios for incorporating the adaptive context factors into our general trust metric. In this section we give two examples of adapting the metric using the transaction and community context factor respectively.

Incorporating Transaction Contexts by Transaction Context Factor

Various transaction contexts, such as the size, category, or time stamp of the transaction and the location information of the transacting peer can be incorporated into the metric. For example, an adapted metric that incorporates the size of a transaction i in terms of the Dollar amount of the payment, denoted by $D(u, i)$, is defined below so the feedback for larger transactions are assigned more weight than those for smaller ones:

$$T(u) = \frac{\sum_{i=1}^{I(u)} S(u,i) * Cr(p(u,i)) * D(u,i)}{\sum_{i=1}^{I(u)} Cr(p(u,i)) * D(u,i)}$$

Providing Incentives to Rate by Community Context Factor

Several remedies have been suggested to the incentive problem of reputation systems such as market-based approaches and policy-based approach in which users will not receive rating information without paying or providing ratings. However, implementing these approaches might stifle the growth of online communities and fledgling electronic markets. In PeerTrust, the incentive problem of reputation systems can be alleviated by building incentives or rewards into the metric through community context factor for peers who provide feedback to others. An adapted metric can be defined below with a reward as a function of the ratio of total number of Feedback peer u give others, denoted as $F(u)$, over the total number of transactions peer u has during the recent time window. The weight factors can be tuned to control the amount of reputation that can be gained by rating others.

$$T(u) = \alpha * \frac{\sum_{i=1}^{I(u)} S(u,i) * Cr(p(u,i))}{\sum_{i=1}^{I(u)} Cr(p(u,i))} + \beta * \frac{F(u)}{I(u)}$$

Evaluation

We performed some initial experiments to evaluate PeerTrust model and show its feasibility, effectiveness, and benefits. The first one evaluates effectiveness of PeerTrust model in terms of its computation error against malicious manipulations of peers in two settings. The second one demonstrates the importance and benefit of supporting reputation based trust in a P2P community by allowing peers to avoid untrustworthy peers using the reputation based trust scheme.

Simulation Setup

Our initial simulated community consists of N peers and N is set to be 128 in most experiments. The game theory research on reputation introduced two types of players

(Dellarocas, 2003). One is commitment type or a long-run player who would always cooperate because cooperation is the action that maximizes the player's lifetime payoffs if the player could credibly commit to an action for the entire duration. In contrast, a strategic type corresponds to an opportunistic player who cheats whenever it is advantageous for him to do. We split peers into these two types in our simulation, namely, good peers and strategic or malicious peers. The percentage of malicious peers is denoted by k . We have one experiment with varying k to show its effect and otherwise k is set to be 25%.

The behavior pattern for good peers is to always cooperate in transactions and provide honest feedback afterwards. While it is a challenging task to model peers' malicious behavior realistically, we start with two malicious behavior patterns to study the robustness of PeerTrust metrics, namely non-collusive setting and collusive setting. In non-collusive setting, malicious peers cheat during transactions and give dishonest ratings to other peers, that is, give bad rating to a peer who cooperates and give good rating to a peer who cheats. A malicious peer may choose to occasionally cooperate in order to confuse other peers and fool the system. We use *mrates* to model the rate that a malicious peer acts maliciously. We have one experiment varying *mrates* to show its effect on trust computation effectiveness, and otherwise *mrates* is set to 100%. In collusive setting, malicious peers act similarly to those in non-collusive setting, and in addition, they form a collusive group and deterministically help each other by performing numerous fake transactions and give good ratings to each other.


We use a binary feedback system where a peer rates the other peer either 0 or 1 according to whether the transaction is satisfactory. The number of transactions each peer has during the latest time window, denoted by I , is set to be 100 for all peers. For comparison purpose, we compare PeerTrust metrics to the conventional approach, referred to as Conventional, in which an average of the ratings is used to measure the trustworthiness of a peer without taking into account the credibility factor. All experiment results are averaged over five runs of the experiments.

Effectiveness against Malicious Behaviors of Peers

The objective of this set of experiments is to evaluate the effectiveness and robustness of the trust metrics against malicious behaviors of peers. The experiments proceeds as peers perform random transactions with each other. After 6,400 transactions in the community, that is, an average of 100 transactions for each peer, a good peer is selected to evaluate the trustworthiness of all other peers. Each experiment is performed under both non-collusive and collusive settings described earlier. We compute the trust computation error as the root-mean-square (RMS) of the computed trust value of all peers and the actual likelihood of peers performing a satisfactory transaction, which is 1 for good peers and $(1 - \text{mrates})$ for malicious peers. A lower RMS indicates a better performance.

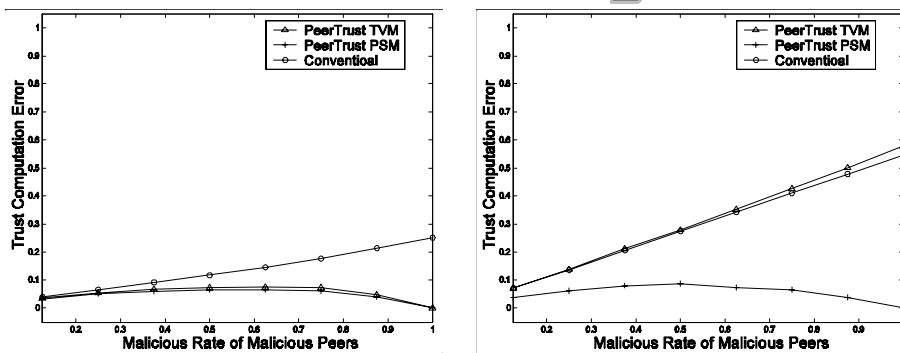
For the first experiment, we vary the percentage of malicious peers (k) and set the malicious rate to 1 ($\text{mrates} = 1$). Figure 1 represents the trust computation error of different PeerTrust algorithms and the conventional approach with respect to k in the two settings.

Figure 1. Trust computation error with respect to percentage of malicious peers in non-collusive setting (top) and collusive setting (bottom)




We can make a number of interesting observations in the non-collusive setting. First, the performance of the conventional approach drops almost linearly when k increases. Without taking into account the credibility of feedback source, it is very sensitive to malicious peers who provide dishonest feedback. Second, PeerTrust TVM stays effective when k is less than 50%. Using trust values of peers recursively as the weight for their feedback, they are able to filter out dishonest feedback and make correct trust computations. However, the error becomes 100% when k is greater than 50%, which indicates they completely make wrong evaluations by mistaking good peers as untrustworthy and malicious peers as trustworthy. This is particularly interesting because it shows that malicious peers are able to fool the system by overriding the honest feedback provided by good peers when they are the majority. Last, PeerTrust PSM stays effective even with a large percentage of malicious peers. This confirms that the personalized similarity based credibility acts as a very effective measure to filter out dishonest feedback. The collusive setting also presents interesting observations. Both conven-

Figure 2. Trust computation error with respect to percentage of malicious peers in non-collusive setting (top) and collusive setting (bottom)



tional metric and PeerTrust TVM metric are extremely sensitive to collusive attempts that dishonestly provide feedback even when the number of malicious peers is very small. On the other hand, PeerTrust PSM metric, as we have expected, acts as a very effective defense against collusion by filtering out dishonest feedback from the collusive group.

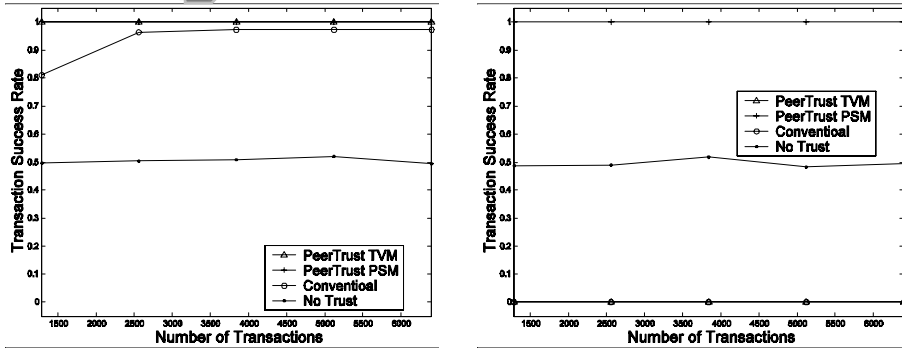
For the second experiment, we vary the malicious rate (*mrate*) and set the percentage of malicious peers to 25% ($k=25\%$). Figure 2 compares the trust computation error of PeerTrust metrics and the conventional metric with respect to *mrate* in the two settings. Again we can make a number of interesting observations in both settings. First, the performance of the conventional approach drops when *mrate* increases. Second, both PeerTrust TVM and PSM metrics have a slightly dropped performance when the malicious rate is less than 100%. This indicates that s are able to confuse the system a little when they occasionally cooperate and give honest feedback. The collusive setting shows similar results but to a larger extent.

Benefit of Trust Based Peer Selection

This set of experiments demonstrates the benefit of using a reputation based trust system in which peers compare the trustworthiness of peers and choose the peer with the highest trust value to interact with. A transaction is considered successful if both of the participating peers cooperate. We define successful transaction rate as the ratio of the number of successful transactions over the total number of transactions in the community up to a certain time. A community with a higher transaction success rate has a higher productivity and a stronger level of security. The experiment proceeds by repeatedly having randomly selected good peers initiating transactions. In a community that has a reputation system, the source peer selects the peer with the highest trust value to perform the transaction. Otherwise it randomly selects a peer. The two peers then perform the transaction and the transaction succeeds only if the selected peer cooperates. The experiment is performed in both non-collusive setting and collusive setting. We show the benefit of utilizing a reputation based trust system that uses conventional and PeerTrust metrics compared to a community without any trust system.

Figure 3 shows the transaction success rate with regard to the number of transactions in the community in the two settings. In the non-collusive setting, we can see an obvious gain of the transaction success rate in communities equipped with a trust mechanism. This confirms that supporting trust is an important feature, as peers are able to avoid untrustworthy peers. We can also see different trust metrics benefit the community to a different extent. This shows a similar comparison to the previous experiment. It is worth noting, however, that the system using conventional metric achieves a transaction success rate close to 100% even though its trust computation error is much higher than 0, shown in Figure 1. This is because even if the computed trust values do not reflect accurately the likelihood of the peers being cooperative, they do differentiate good peers from bad peers in most cases by the relative ranking. In the collusive setting, we can see that the transaction success rate is 0 for the system using conventional and PeerTrust TVM metric. This indicates that malicious peers are able to completely fool these trust schemes by collusion and render the system useless, even worse than the system without

Figure 3. Benefit of reputation based trust scheme in non-collusive (top) and collusive setting (bottom)



a trust scheme. However, the system still benefits from PeerTrust PSM metric significantly and shows robustness against the collusion.

Conclusion and Future Trends

We discussed reputation and trust and described PeerTrust model for building reputation based trust systems for e-commerce including m-commerce applications. It alleviates or avoids some of the security risks we discussed earlier by helping participants to choose reputable participants and avoid untrustworthy ones. For example, the simplest version of a virus attack would be that an adversary delivers a virus to a good peer or member. With a reputation based trust mechanism in place, the peer who receives the malicious content will be able to submit a negative feedback about the malicious peer and help other peers to avoid it in the future.

Not surprisingly, a reputation-based trust mechanism also introduces vulnerabilities and problems by itself. Common attacks are known as shilling attacks, where adversaries attack the system by submitting fake or misleading ratings to confuse the system as we have discussed earlier. Further, participants can amount attacks on the trust management system by distributing tampered with trust information. PeerTrust tries to minimize such security weaknesses. For example, the use of the credibility factor of the feedback source can be seen as an effective step towards handling fake or misleading ratings. Reputation-based feedback. The ability to incorporate various transaction and community contexts can also act against some of the subtle attacks. Furthermore, by combining the proposed trust metric and the secure trust data transmission built on top of public key cryptographic algorithms, it prevents distribution of tampered with trust information and man in the middle attack.

There remain many interesting research problems, some of which are listed below:

- *Collusion among participants.* Unfortunately there is so far no mechanism that can completely prevent this type of attack. Developing mechanisms that are robust to collusion among participants is currently an active area for research.
- *Lack of portability of reputation between systems.* This limits the effectiveness of reputation systems. For example, if a mobile user travels to a foreign network, he or she would become a newcomer in that network and lose all his/her reputation in his or her home network. Efforts are currently underway to construct a more universal framework in e-commerce research. However, it is yet to receive a global acceptance.
- *Get rid of bad history through reentry.* Another risk mainly in the P2P community is that peers can easily discard their old identity and adopt a new one through reentry to get rid of the bad history. Potentially there are two classes of approaches to this issue: either make it more difficult to change online identities, or structure the community in such a way that exit and reentry with a new identity becomes unprofitable (Friedman, 2001).
- *One-time attack.* The proposed trust building techniques are based on experiences. Therefore, a peer that has been consistently reliable can perform an unavoidable one-time attack. Although trust metrics can be adapted to quickly detect a malicious participant's bad behavior, it is very hard if not impossible to fully prevent this type of attack.

We believe efforts for promoting reputation and trust play an important role in m-commerce security, which is a key to the acceptance and general deployment of m-commerce applications.

Acknowledgment

We would like to thank the reviewer of this chapter and the editors of the book for their valuable comments.



References

- Chari, S., Kerani, P., Smith, S., & Tassiulas, L. (2001). Security issues in m-commerce: A usage-based taxonomy. *E-Commerce Agents, LNAI 2033*, 264-282.
- Dellarocas, C. (2001). *Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms*. 3rd ACM Conference on Electronic Commerce.

- Dellarocas, C. (2003). The digitization of word-of-mouth: Promise and challenges of online reputation mechanism. *Management Science*, 49(10), 1407-1424.
- Friedman, E., & Resnick, P. (2001). The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(1), 173-199.
- Manchala, D.W. (2000). E-commerce trust metrics and models. *IEEE Internet Computing*, 4(2), 36-44.
- Resnick, P., Zeckhauser, R., Friedman, E., & Kuwabara, K. (2000). Reputation systems. *Communications of the ACM*, 43(12), 45-48.
- Varshney, U., & Vetter, R. (2002). Mobile commerce: Framework, applications and networking support. *ACM/Kluwer Journal on Mobile Networks and Applications*, 7(3), 185-198.
- Xiong, L., & Liu, L. (2003). *A reputation-based trust model for peer-to-peer eCommerce communities*. IEEE Conference on Electronic Commerce.

Endnotes

- ¹ PalmOS/Phage.963 virus. http://vil.nai.com/vil/content/v_98836.htm
- ² PalmOS/LibertyCrack virus. http://vil.nai.com/vil/content/v_98801.htm
- ³ Funny SMS Messages. <http://www.free-sms-messages.com/viruses.html>
- ⁴ Wireless Application Protocol Forum. WAP Wireless Transport Layer Security. Version 06-Apr-2001.
- ⁵ Wireless Application Protocol Forum. WAP Public Key Infrastructure. Version 24-Apr-2001.