Regular Paper

A Reputation Management Scheme for Peer-to-Peer Networks based on the EigenTrust Trust Management Algorithm

Takuya Nishikawa[1,a) Satoshi Fujita[1]

Received: September 7, 2011, Accepted: December 16, 2011

Abstract: Peer-to-Peer (P2P) systems have attracted considerable attention in recent years, as a key technology to realize scalable, dependable network services. However, because of its high anonymity, P2P systems involve several drawbacks such as the weakness against malicious attacks by anonymous peers. In this paper, we propose a method to evaluate the trustworthiness of each peer by explicitly taking into account the accuracy of mutual evaluations. The proposed method is an extension of the EigenTrust proposed by Kamvar et al. which calculates a global trust vector consistent with the observed local trust vectors under the weighted sum in a linear space. The performance of the proposed method is evaluated by simulation. The result of simulations indicates that the proposed method identifies a large subset of reliable peers with a sufficiently small number of message transmissions compared with a simple modification of the EigenTrust.

Keywords: peer-to-peer, reputation, trust

1. Introduction

Trust management is a key issue in realizing dependable computer systems in an open environment such as the Internet. Thus far, many trust management schemes have been proposed in the literature [3], [6], [7], [8], [11], [14], [15]. The main objective of trust management is to make assessments and decisions regarding the dependability of potential transactions involving risk, and to allow users and the system owners to increase and correctly represent the reliability of themselves and their systems [5]. In particular, the goal of a trust management in distributed environments is to identify potential risk of operations conducted by malicious users and to certify the dependability of distributed systems using several sources of information, such as the access log of each user, stochastic behavior of the users, and the reputation about the reliability of the other users.

Among them, the "reputation" of users is considered to be a rich source of dependability information for many online applications, to assess the reliability of each user in a distributed environment. In fact, many online auction systems such as eBay [*1], eBid [*2], and OZtion [*3] try to increase their dependability by allowing customers to make an assessment of their past counterparty, and by disclosing the result of such assessments to all users. By referring to those assessments, each user can identify malicious users who involve potential risks, such as the possibility of the download of inauthentic files, a long delay of transactions and a sudden cancellation of ongoing transactions.

A critical problem in such a reputation-based trust management is that the reputation given by a malicious user may not be reliable. Although it would be possible to omit "all" reputations given by suspected users, we cannot identify a sufficient number of reliable users under such a pessimistic approach. We need to carefully take into account the reputation given by every user (including suspected ones) to identify as large number of reliable users as possible. In this paper, we propose a systematic way to realize such an assessment. As a basic infrastructure, we focus on the EigenTrust trust management system [8] proposed by Kamvar et al. As will be described later, the main idea of the EigenTrust is to deal with the transitivity of the trust of users in a linear space; i.e., under the model used in the EigenTrust, if a evaluates the reliability of b as 0.5, and b evaluates the reliability of c as 0.5, then a (indirectly) evaluates the reliability of c as $0.5 \times 0.5 = 0.25$. Such an approach would work well if every user can make an assessment accurately, but in actual situations, the assessment conducted by unreliable users may not be accurate, and such an inaccuracy increases the risk of overlooking malicious users in identifying a set of reliable users.

In the proposed reputation management scheme, we take into account such an inaccurate evaluation by explicitly assigning low priority to those evaluations. More concretely, we propose two different ways of selecting reliable neighbors during the calculation of the trust values under the EigenTrust algorithm. The accuracy and the efficiency of the proposed scheme are evaluated by simulation. In addition to the previous study [1], we conducted some additional simulations to evaluate the effectiveness of the

[1] Hiroshima University, HigashiHiroshima, Hiroshima 739–8527, Japan
[a)] taku42@se.hiroshima-u.ac.jp

[*1] http://www.ebay.com/

[*2] http://www.ebid.net/

[*3] http://www.oztion.com.au/

proposed scheme in more detail. The result of simulations indicates that the proposed scheme can improve the accuracy of the original EigenTrust, and simultaneously, can significantly reduce the amount of message transmissions.

The remainder of this paper is organized as follows. Section 2 describes the EigenTrust algorithm and its variants. Section 3 describes the proposed method. The result of simulations is summarized in Section 4. Finally, Section 5 concludes the paper with topics for future works.

2. EigenTrust

Consider a P2P system consisting of N homogeneous peers. Each peer i is assigned a real number $t_i^* \in [0, 1]$ called **credibility**, where each peer is regarded as being reliable if it is assigned credibility close to one. The credibility of a peer can not be directly referred to by the other peers, while we assume that the credibility of peer j can be "evaluated" by any peer in the system by conducting a transaction with j, and will say that "peer i trusts j" if i evaluates that j is reliable.

2.1 Algorithm

EigenTrust [8] is a reputation management algorithm based on the notion of *transitive trust*, in the sense that if peer i trusts peer j, then peer i would also trust peers trusted by j. Let c_{ij} denote the **local trust value** of j which is locally calculated by i through transactions with j such as a file download and a query processing. More concretely, it is calculated by the number of satisfactory or unsatisfactory transactions conducted between i and j, and is normalized to a real number in $[0, 1]$ (if there were no transactions between them, c_{ij} is calculated to be 0). Given such a set of local trust values, peer i can *estimate* the credibility of peer k, in the following manner:

$$t_{ik} = \sum_j c_{ij} \cdot c_{jk}, \qquad (1)$$

that is, the estimated credibility t_{ik} is calculated by collecting all local trust values for peer k, and by taking a weighted sum of those values, where the weight of value c_{ij} is the local trust value of j calculated by i. Let $\vec{c_i}$ denote the local trust vector calculated by i, and $\vec{t_i}$ denote a vector of estimated credibility calculated by using Eq. (1). Then, by using **reputation matrix** C, which consist of local trust values collected from the other peers, Eq. (1) can be restated as follows:

$$\vec{t_i} = C^T \cdot \vec{c_i}. \qquad (2)$$

The reader should note that in the above equation, $\vec{t_i}$ reflects the local trust values calculated by the immediate neighbors of i, whereas $\vec{c_i}$ merely reflects the local trust values of i. This indicates that, by multiplying C^T to $\vec{t_i}$ from the left, we can obtain a refined trust vector reflecting the local trust values of the peers within distance of two from i. By repeating similar operations, we will have a stationary vector $\vec{t}$ such that

$$\vec{t} = (C^T)^x \cdot \vec{c_i} \qquad (3)$$

for some $x \geq 1$, since if C is aperiodic and strongly connected, then powers of reputation matrix C converges to a stable value at some point. Such a vector $\vec{t}$ is known as the left principal eigenvector of the reputation matrix C, and since it is independent of the selection of peer i, it can be regarded as a global trust vector in the system, in the sense that the j^{th} element in $\vec{t}$ represents the **global trust value** of peer j. It would be worth noting here that the resulting vector $\vec{t}$ may *not* coincide with the real credibility vector, since the above observation merely claims that the vector consistently explains the local trust values under the "weighted sum" represented in Eq. (1).

The calculation of such a stationary vector $\vec{t}$ can be conducted in a straightforward and approximated manner. More concretely, by assuming that each peer knows the whole matrix C, it may repeatedly left multiply the current vector by C^T, until the difference to the previous vector becomes smaller than a predetermined threshold.

2.2 Variants

In the literature, several variants of the EigenTrust have been proposed.

Abrams et al. [2] introduced the notion of *cyclic partitioning* to the EigenTrust, to prevent selfish peers from maliciously increasing their credibility by manipulating the local trust values of the colluding peers. In this method, the set of peers are evenly partitioned into m subsets (i.e., colors), and those colors are arranged into a directed cycle. Each color has $\lfloor \frac{N}{m} \rfloor$ or $\lceil \frac{N}{m} \rceil$ peers, where N denotes the number of peers in the P2P system. Each peer in color c can send a query and receive a service merely from a peer in its successor color $succ(c)$ in the directed cycle, but cannot directly evaluate the peers in $succ(c)$. More concretely, the outgoing links from color c are set to be uniform over $succ(c)$, and then during the calculation of the global trust vector for peers in a particular color c, we should replace element c_{ij} in the reputation matrix C by m/N if $i \in c$. By this modification, we can realize a situation in which the local trust values provided by each peer in color c do not directly affect the calculation of the global trust vector of the peers in color c, and it effectively removes an incentive to give a wrong reputation to itself or the colluding peers.

Donato et al. [4] assume the existence of two different types of malicious peers, and proposed several extensions of the EigenTrust. The first type is a peer which always uploads inauthentic files. Peers of the second type, called *spies*, return authentic files only to popular queries, and return inauthentic files to the other queries. Both types of peers assign positive local trust values only to malicious peers.

In order to eliminate the effect of such malicious peers, Donato et al. introduced a logical network called positive opinion network in which a directed link is inserted from peer i to peer j only when i downloads authentic files from j, where each link is weighted by the number of authentic files downloaded from j. *Inverse EigenTrust* is a reputation management scheme which applies the EigenTrust to the transpose of a positive opinion network called an inverse network. By definition, no good peer can reach a malicious peer through a directed path in an inverse network. Therefore, under the Inverse EigenTrust, a good peer can always calculate the score of malicious peers to be zero. Let $I(i)$ be the global trust value of i calculated under the Inverse Eigen-

Trust, and $E(i)$ be the global trust value calculated by applying the original EigenTrust to the positive opinion network. In the Donato et al.'s scheme, in order to disregard the opinion of spies, the global trust value of peer i is calculated to be 0 if $I(i) = 0$, otherwise it is calculated to be $E(i)$. Note that the first condition can be relaxed to $I(i) < \theta$ for some threshold θ, in order to eliminate cases in which malicious peers assign positive local trust values to a small percentage of good peers, since peers can reach malicious peers on an inverse network and in such a situation, $I(i)$ of malicious peers i cannot be 0.

3. Main Results

3.1 Overview

In this paper, we extend the reputation management conducted under the EigenTrust in the following three directions. In the first direction, we will introduce the notion of *certainty* of local trust values calculated by each participating peer. Recall that in the original EigenTrust, such a certainty was partially taken into account as a weight of the peer while conducting a summation of the elements contained in the current vector $\vec{t_i}$, and it is based on an assumption that if peer j is evaluated to be unreliable by a peer, then it would also be evaluated to be unreliable by any other peers. However, such a simple method does not exclude the effect of a malicious behavior of unreliable peers such that j behaves honestly during transactions with i but it behaves dishonestly during transactions with i'. In the following, we model such an uncertainty by using a probabilistic approach, and extend the EigenTrust to take into account such malicious behaviors of unreliable peers.

The second direction is a reduction of the number of message transmissions during the calculation of a global trust vector in a distributed environment. As was described earlier, the EigenTrust assumes that every peer knows the whole reputation matrix before starting the calculation, but it requires a large amount of message transmissions, while it tries to reduce such traffic by using a distributed data structure such as Distributed Hash Table (DHT) [10], [12], [13]. Thus, if we could effectively skip the communication to a number of (unreliable) peers, it could significantly reduce the traffic of the scheme.

The third direction is to avoid sacrificing the quality of the solution. If we skip the communication to several peers, the amount of local trust values we can collect will decrease and the accuracy of calculation of global trust vectors may become low. In order to overcome such an inaccuracy of calculation, we introduce a method which enables us to collect local trust values from a sufficient number of peers without significantly increasing the amount of message transmissions.

The modifications to the EigenTrust are summarized in **Table 1**. The details of the modifications are described in the following sections.

3.2 Model of Unreliable Peers

In the proposed method, we assume that an evaluation conducted by a reliable peer is accurate, but an evaluation conducted by an unreliable peer may not be accurate. In order to formalize such an uncertainty, we introduce a **range-based approximation model** described below (in what follows, we call this model RBA); when i evaluates j, the local trust value c_{ij} is randomly selected from the following range with uniform probability:

Table 1 Difference of the proposed scheme (CRD stands for credibility).

	proposed scheme	EigenTrust
Certainty of CRD	Blurred	Not blurred
Collection of CRD	Within fixed TTL	Whole network
Spreading of CRD	Yes	No
Pre-trusted peers	Not exist	Exist

$$\left[\max\{0, t_j^* + t_i^* - 1\}, \min\{1, t_j^* - t_i^* + 1\}\right].$$

By definition, the resulting value c_{ij} coincides with the credibility t_j^* of j if $t_i^* = 1$, and the value will be "blurred" as decreasing the value of t_i^*; e.g., if $t_i^* = 0.8$, such a blur increases to 0.2 $(= 1 - t_i^*)$, and the value of c_{ij} becomes a random real number in $[\max\{0, t_j^* - 0.2\}, \min\{1, t_j^* + 0.2\}]$.

In addition to the above assumption, in the following, we will assume that there are no absolutely reliable peers whose existence is known to all peers. Such peers are called "pre-trusted" peers in Ref. [8]. Pre-trusted peers play an important role in improving the efficiency in the original EigenTrust algorithm, although in practice, it is hard to assume the existence of such peers.

3.3 Problem

In this paper, we consider a problem of identifying reliable peers from a collection of local trust values, which is informally stated as follows: *Given a set of local trust values distributed over a network, identify a set of reliable peers as accurate as possible with as small number of message transmissions as possible*. Here, we should notice that in many practical situations, the risk of mis-identifying unreliable peers as a reliable peer is much higher than the risk of mis-identifying reliable peers as an unreliable peer. In other words, a concrete goal of the above abstract problem should be to identify a *large subset of reliable peers* with as small number of message transmissions as possible. In the following, we call this problem MAX_REL. Note that it is completely different from the reputation management problem considered in the EigenTrust.

In fact, it is easy to see that the global trust vector calculated by the EigenTrust is not sufficient as a solution to the problem MAX_REL under the RBA model. To see this in detail, let us consider an example consisting of three peers i, j, and k. Suppose $t_i^* = 1$, $c_{ij} = 0.7$ and $c_{jk} = 0.5$. In the EigenTrust algorithm, peer i estimates the credibility of k through j as

$$c_{ij} \times c_{jk} = 0.7 \times 0.5 = 0.35.$$

On the other hand, since $t_i^* = 1$, c_{ij} coincides with the credibility of j, i.e., $t_j^* = 0.7$. Thus, the credibility of peer k is at least 0.2 $(= 0.5 - 0.3)$ and at most 0.8 $(= 0.5 + 0.3)$; i.e., there is risk in believing that the credibility of k is 0.35 since there is a possibility that the credibility of the peer is 0.2.

3.4 Algorithm

Under the RBA model, the local trust value calculated by each unreliable peer is not accurate in the sense that the degree of inaccuracy is proportional to the credibility of the evaluator. This means that the calculation result of unreliable peers involves a risk of an unexpected situation in which *an unreliable peer is mis-identified to be reliable*. In order to avoid such situations, in the proposed scheme, we will take an approach to selectively omit the local trust value calculated by unreliable peers. It is worth noting here that to realize such an omission in an effective manner, we have to make a decision on the reliability of particular peers before completing the calculation of their credibility. In addition, we should avoid omitting all suspected values, since our objective is to identify as large subset of reliable peers as possible.

3.4.1 Selective Query Forwarding

The basic idea of the proposed scheme is as follows. Consider a directed graph G such that the vertex set is the set of peers, and for any pair of peers i and j, an edge is placed from i to j iff $c_{ij} > 0$. Peer i tries to collect local trust values from peers in the graph by disseminating a query message along directed edges in G, in such a way that: 1) the query reaches peers which have high reliability, and 2) the query does not reach peer j if it is not connected with i through a path consisting of reliable peers. A query has a fixed TTL (Time To Live), and the range of the region to which the query reaches is restricted by the TTL. Each peer who received a query from an adjacent peer returns its local trust values to the source of the query message.

More concretely, in the proposed method, we adopt the following two query forwarding rules:

- In the first rule, peer i forwards the received query to all neighbors j such that $c_{ij} \geq \theta_i$, where θ_i is the average of local trust values calculated by i.
- In the second rule, peer i forwards the received query to its neighbors in a non-increasing order of their local trust value, until the sum of local trust values of the peer who have received the query exceeds a threshold θ'_i. The threshold θ'_i is determined as the percentage of the number of neighbors such that $c_{ij} \geq \theta_i$, where θ_i is the threshold used in the first rule.

In both rules, peer i does not forward a query if it is received from k such that $c_{ik} = \max_j\{c_{ij}\}$, in order to avoid unnecessary query forwarding to the peers with low trust values. We call a neighbor of i as *Good* for i if it receives a query from i. Note that in each rule, each peer forwards a query only to its Good neighbors.

3.4.2 Spreading of Local Trust Values

In addition to the above query forwarding rules, in the proposed scheme, we adopt the following spreading scheme which will be used with the selective query forwarding scheme in a combined manner. The aim of the spreading scheme is to disseminate the local trust value of each peer to its neighbors beforehand, to reduce the cost of query forwarding. The spread of local trust values is controlled by a TTL. A local trust vector received from a neighbor is cached into its local storage, and is forwarded to other neighbors only if it comes from a Good neighbor. Then, the overall query forwarding scheme is modified in such a way that each peer returns a collection of local trust values cached in its local storage (including its own local trust values) as a response to a query received from an adjacent peer.

Local trust values which could not be obtained through the above schemes are simply assumed to be zero, and each peer locally applies the EigenTrust algorithm to the collected partial information to calculate an approximated global trust vector. Each peer then determines the reliability of the other peers by referring to the calculated global trust values, e.g., it can output a subset of peers whose global trust value exceeds a threshold as a set of reliable peers. The reader should note that, since each peer keeps its own reputation matrix, each peer obtains unique global trust vectors through the proposed scheme although every peer obtains the same global trust vector in the original EigenTrust algorithm.

4. Experiments

In this section, we compare the effectiveness of the proposed scheme with conventional reputation management schemes with respect to the accuracy of estimation and the number of message transmissions.

4.1 Metric

Under the RBA model, the local trust value calculated by each peer involves an inaccuracy proportional to the credibility of the evaluator, and it takes a random value in a range determined by the inaccuracy. This implies that it is impossible to exactly estimate the actual credibility by any scheme under the model. In the following, we will evaluate the accuracy of reputation management schemes by the number of mis-identifications under the model. Recall that in each reputation management scheme considered in this paper, each peer i autonomously conducts a calculation of the global trust vector, and the resulting vector may differ for each i.

Let X denote a list of peers which is arranged in a non-increasing order of credibility. Let $\vec{t_i}$ denote the global trust vector calculated by i, and let Y_i denote a list of peers which is arranged in a non-increasing order of the values in vector $\vec{t_i}$. For example, if the global trust vector $\vec{t_i}$ is calculated as

$$(0.3 \ \ 0.5 \ \ 0.2 \ \ 0.9)$$

then the list Y_i is determined as

$$(4, 2, 1, 3)$$

since value 0.9 for peer 4 is the largest, value 0.5 for peer 2 is the second largest, and so on.

As was mentioned earlier, the risk of mis-identifying unreliable peers as a reliable peer is much higher than the risk of mis-identifying reliable peers as an unreliable peer. Thus in the following, we evaluate the risk of the resulting list Y_i, by counting the number of peers in the bottom σ% of list X (i.e., peers whose credibility is low) which appear in the top σ% of list Y_i (i.e., peers mis-identified to be reliable by the scheme). We will analyze the accuracy of the schemes by setting σ to 30%. The number of such mis-identifications shall be bound to as small a number as possible. In the following, for brevity, we call a peer with a low credibility a *incredible peer*, and a peer with a high credibility a *credible peer*.

Table 2 summarizes parameters fixed in the simulations. In the following, the other parameters such as LTP, TTL_q, TTL_s, and σ will be determined appropriately, where LTP denotes the percentage of incredible peers in the network, TTL_q denotes the TTL

Table 2 Simulation settings.

# of peers	1000
average # of neighbors (degree)	6.0
minimum # of neighbors (degree)	3
credibility of credible peers	[0.8, 1.0]
credibility of incredible peers	[0.0, 0.2]

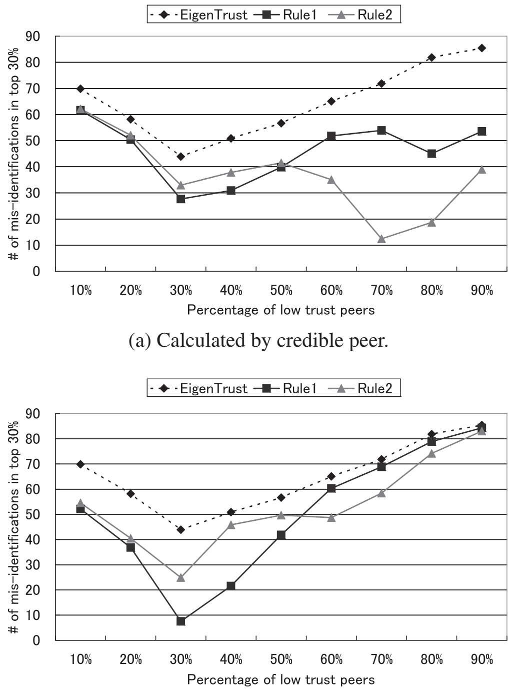

(a) Calculated by credible peer.

(b) Calculated by incredible peer.

Fig. 1 The number of mis-identifications in the top 30% of Y_i.

used in the query forwarding scheme and TTL_s denotes the TTL used in the spreading scheme.

4.2 Comparison with EigenTrust

Figure 1 illustrates the accuracy of the proposed method. The horizontal axis of the figure is the LTP which varies from 10% to 90%, and the vertical axis is the percentage of mis-identifications, where other parameters are fixed as $\mathsf{TTL}_q = 3$ and $\mathsf{TTL}_s = 2$. Figure 1 (a) shows the accuracy of the estimation calculated by credible peers, and Fig. 1 (b) shows that of incredible peers. We can observe from the figures that for any value of the LTP, the proposed peer selection rules certainly exhibit a better performance than the EigenTrust algorithm, and in the calculation conducted by credible peers, the second rule is particularly effective to reduce the number of mis-identifications. In Fig. 1 (b), however, when the percentage of incredible peers exceeds 30%, the number of mis-identifications increases rapidly. This is because incredible peers are likely to estimate neighbors inaccurately under the RBA model, and the behavior of peers such that a query is forwarded only to Good neighbors does not work well. It would be noted that although we could observe another increase of mis-identifications for small LTPs such as 10% or 20% (from right to left), there is no need to worry about. In fact, when LTP=10%, the upper 20% of the bottom 30% of the list X should consist of credible peers (recall that LTP=10% implies that 90% of the peers are credible peers). Thus, even if all of such peers are estimated as credible peers, it does not cause a risk.

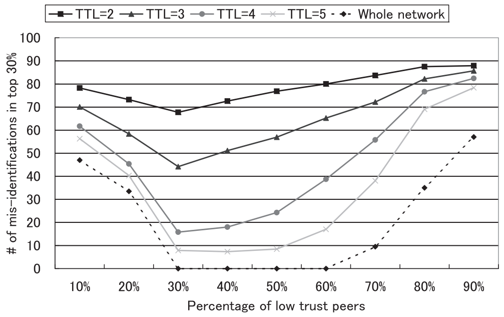

Fig. 2 The impact of TTL_q in the original EigenTrust.

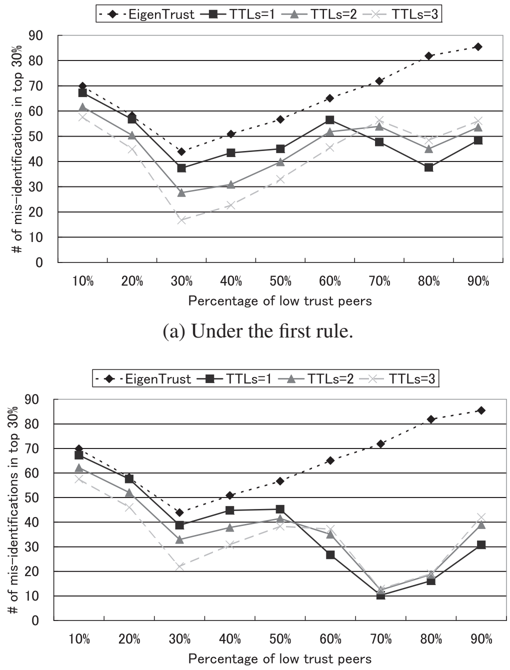

(a) Under the first rule.

(b) Under the second rule.

Fig. 3 The impact of TTL_s.

Figure 2 shows the accuracy of the estimation calculated by the EigenTrust algorithm when the range of queries is restricted by TTL_q. We varied TTL_q from 2 to 5, or set to ∞. We can observe from the figure that the number of mis-identifications increases as decreasing TTL_q, i.e., the restriction of the range of query forwarding by a simple TTL is not sufficient under the RBA model, while it is effective to some extent in reduction in the amount of the traffic. In addition, even if TTL_q is set to ∞, the number of mis-identifications rapidly increases when LTP exceeds 60%. This suggests that selective query forwarding would be necessary to efficiently improve the accuracy of the estimation.

Figure 3 illustrates the impact of TTL_s to the accuracy of the schemes. As shown in the figure, for small LTPs, the accuracy increases with an increasing TTL_s, since each peer can collect a large amount of local trust values from credible peers. In addition, for large LTPs, the number of mis-identifications rapidly decreases, particularly under the second rule, although such a tendency becomes small as the TTL_s increases.

4.3 Analysis

In our proposed schemes, each peer keeps its own reputation matrix C. In addition, an entry corresponding to peer j in the

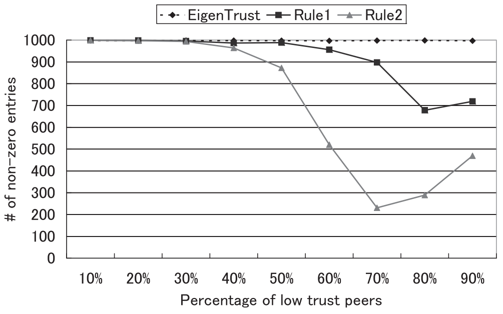

Fig. 4 The number of non-zero entries.

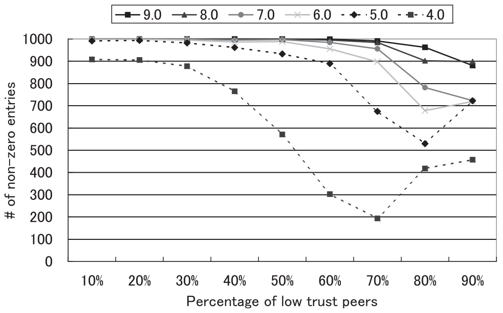

Fig. 5 The impact of average degree against non-zero entries.

global trust vector calculated by peer *i* from its reputation matrix *C* takes value zero if *i* is not connected to *j* through a directed path consisting of reliable peers. In a real world P2P system, this situation implies that peer *i* might have to interact with peer *j* without considering whether it is reliable. Such an "invisibility" of unreliable peers plays a key role in the proposed schemes to reduce the risk of unreliable reputations. To verify such conjecture, we counted the number of (visible) peers whose trust values are non-zero in a calculated global trust vector. **Figure 4** illustrates the result. As shown in the figure, for large LTPs, queries issued in the proposed schemes certainly reach a very restricted region of the network, so that the amount of non-zero entries becomes low. For example, when LTP=70%, the percentage of non-zero entries is 90% under the first rule, and it is only 25% under the second rule. This is why the number of mis-identifications sharply reduces under the second rule in Fig. 3 (b).

Although such a reduction of non-zero entries plays a crucial role in the proposed schemes, we should avoid a disappearance of such entries to effectively identify a set of reliable peers. We could keep sufficient number of non-zero entries, if each query reaches a peer with many adjacent peers. To verify this conjecture, we evaluate the number of non-zero entries by varying the average degree of peers from 4.0 to 9.0. **Figure 5** shows the results. As the average degree increases, the number of non-zero entries also increases even for large LTP's. In fact, a high degree peer effectively disseminates queries, and has many local trust values by itself. Thus, by acquiring those values from high degree peers, each peer can locally have a dense reputation matrix *C*, which increases the possibility of calculating many non-zero entries.

The result of the above experiments could be summarized as follows: 1) the first rule slightly beats the accuracy of the second rule for small LTPs particularly when such a calculation is conducted by a incredible peer (see Fig. 1 (b)), while 2) the second rule beats the first one for large LTPs (see Fig. 1 (a)). A remarkable advantage of the first rule is that the number of non-zero entries is not too small compared with the second rule. In fact, under the second rule, the number of non-zero entries becomes very small so that we could not estimate the global trust value of many peers for large LTP's. On the other hand, under the first rule, the number of non-zero entries is kept to be sufficiently large for any value of LTP.

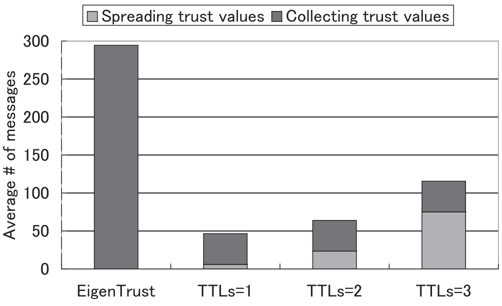

Fig. 6 The number of message transmissions.

4.4 Message Transmissions

Finally, we evaluate the number of message transmissions of the proposed schemes. In the evaluation, we accumulated messages issued for spreading local trust values and queries for collecting local trust values. **Figure 6** summarizes the result. As shown in the figure, the proposed scheme significantly reduce the amount of messages, although it gradually increases along with an increasing the TTL used in the spreading scheme.

In summary, we can conclude that the proposed scheme can improve the accuracy of the original EigenTrust even in situations where many malicious peers randomly report inaccurate local trust values. Moreover, the proposed schemes can significantly reduce the amount of message transmissions.

5. Concluding Remarks

In this paper, we proposed a model of the trust of peers which reflects the inaccuracy of evaluation conducted by unreliable peers. We then proposed an improvement of the EigenTrust algorithm which increases the accuracy of estimation while reducing the traffic of the underlying P2P overlay. The performance of the proposed schemes is evaluated by simulations.

An interesting direction for further research is to compare the performance of the proposed scheme with other trust models and schemes.

References

[1] Nishikawa, T. and Fujita, S.: An Effective Risk Avoidance Scheme for the EigenTrust Reputation Management System, *2010 1st International Conference on Networking and Computing*, pp.36–43 (2010).

[2] Abrams, Z., McGrew, R. and Plotkin, S.: A non-manipurable trust system based on EigenTrust, *SIGecom Exchange*, Vol.5, No.4, pp.21–30 (2005).

[3] Damiani E., Vimercati, D.C. and Paraboschi, S.: A reputation-based approach for choosing reliable resources in peer-to-peer networks, *Proc. 9th ACM Conference on Computer and Communications Security*, pp.207–216 (2002).

[4] Donato, D., Paniccia, M., Selis, M., Castillo, C., Cortese, G. and Leonardi, S.: New metrics for reputation management in P2P networks, *International Workshop on Adversarial Information Retrieval on the Web*, pp.65–72 (2007).

[5] Grudzewski, W.M., Hejduk, I.K., Sankowska, A. and Watuchowicz, M.: *Trust Management in Virtual Work Environments: A Human Factors Perspective*, CRC Press Taylor & Francis Group, p.38 (2008).

[6] Ito, Y. and Kawano, H.: Evaluation of P2P contents distribution system with cryptographic trust chains, *DBSJ Letters*, Vol.6, No.1, pp.21–24 (2007).

[7] Jin, Y., Zhang, Y., Qu, W., Liu, Y. and Li, K.: A trust model based on similarity evaluation in P2P networks, *Proc. International Symposium on Parallel and Distributed Procesing with Applications*, pp.737–742 (2008).

[8] Kamvar, S.D., Schlosser, M.T. and Garcia-Molina, H.: The EigenTrust algorithm for reputation management in P2P networks, *Proc. 12th International Conference on World Wide Web*, pp.640–651 (2003).

[9] Marti, S. and Garcia-Molina, H.: Taxonomy of trust: Categorizing P2P reputation systems, *Computer Networks*, Vol.50, No.4, pp.472–484 (2006).

[10] Maymounkov, P. and Mazières, D.: Kademlia: A peer-to-peer information system based on the XOR metric, *Proc. International Workshop on Peer-to-Peer Systems*, pp.53–65 (2002).

[11] Ooi, B.C., Liau, C.Y. and Tan, K.L.: Managing trust in peer-to-peer systems using reputation-based techniques, *Proc. 4th International Conference in Advances in Web-age Information Management*, pp.17–19 (2003).

[12] Ratnasamy, S., Francis, P., Handley, M., Karp, R. and Shenker, S.: A scalable content-addressable network, *Proc. ACM SIGCOMM*, pp.161–172 (2001).

[13] Stoica, I., Morris, R., Liben-Nowell, D., Karger, D., Kaashoek, M.F., Dabek, F. and Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications, *Proc. ACM SIGCOMM*, pp.149–160 (2001).

[14] Wang, Y. and Vassileva, J.: Trust and reputation model in peer-to-peer networks, *Proc. 3rd International Conference on Peer-to-Peer Computing*, pp.150–157 (2003).

[15] Zhou, R. and Hwang, K.: PowerTrust: A robust and scalable reputation system for trusted peer-to-peer computing, *IEEE Trans. Parallel and Distributed Systems*, Vol.18, No.4, pp.460–473 (2007).

Takuya Nishikawa received his B.E. degree in information engineering from Hiroshima University in 2010. His current research interests include peer-to-peer networks and distributed systems.

Satoshi Fujita received his B.E. degree in electrical engineering, M.E. degree in systems engineering, and Dr.E. degree in information engineering from Hiroshima University in 1985, 1987, and 1990, respectively. He is a Professor at Faculty of Engineering, Hiroshima University. His research interests include communication algorithms on interconnection networks, parallel algorithms, graph algorithms, and parallel and distributed computer systems. He is a member of IPSJ, SIAM Japan, IEEE, and SIAM.