



Trusted P2P Transactions with Fuzzy Reputation Aggregation

Internet commerce and online commodity exchanges suffer from distrust among sellers and buyers, who are often strangers to each other. The authors present a new P2P reputation system based on fuzzy logic inferences, which can better handle uncertainty, fuzziness, and incomplete information in peer trust reports. This system aggregates peer reputations with affordable message overhead. By testing the system using eBay transaction data in the public domain, the authors seek to demonstrate the efficacy and robustness of two P2P reputation systems — FuzzyTrust and EigenTrust — at establishing trust among the peers in P2P applications.

Shanshan Song, Kai Hwang, and Runfang Zhou
University of Southern California

Yu-Kwong Kwok
University of Hong Kong

The Internet has enabled e-commerce and e-auctions for online commodity exchanges among strangers worldwide. Many companies, including eBay, Amazon, uBid, and Yahoo, now offer e-auction platforms to facilitate such transactions. A growing number of online commercial transactions occur in peer-to-peer (P2P) environments, which don't require a central authority to mediate such exchanges. Rather, participating peers can sign in and out asynchronously at will and perform their transactions point-to-point or point-to-multipoint anonymously. Thus, a business or enterprise must have an effective reputation system to help users locate trustworthy partners and exchange commodities securely with confidence.

A reputation system's effectiveness depends on the trust model adapted. Most existing trust models are developed for general-purpose P2P applications that emphasize distributed file-sharing¹ (see the "Related Work in P2P Reputation Systems" sidebar). Little work has been done in establishing trust in e-commerce applications. Most notably, Daniel Manchala has assessed some metrics to quantify the trust among e-transaction intermediaries.² Li Xiong and Ling Liu proposed the PeerTrust system for trust management.³

At the University of Southern California (USC), we aimed to develop an effective and efficient reputation system based on a fuzzy-logic approach, leveraging fuzzy-logic's ability to handle uncertain-

Related Work in P2P Trust Management

Several representative P2P reputation systems currently exist, although the list we present is by no means exhaustive. The eBay and PeerTrust systems introduced here focus on trust management in securing commodity exchanges in e-commerce applications, as does the FuzzyTrust system we present in the main text. The other systems focus on generic P2P applications such as P2P file sharing and Web service sharing.

The eBay user feedback system¹ (www.ebay.com) is by far the simplest and most popular trust-management system, and is specifically tailored for e-auction applications. It applies a centralized database to store and manage the trust scores. Data is open to the general public, so newcomers can easily obtain peer scores. It's really a hybrid P2P system using both distributed client resources and centralized servers.

R. Guha and colleagues at IBM Almaden Research Center developed an interesting idea about the propagation of distrust.² In addition to maintaining positive trust values for peers, the system also allows the proactive dissemination of some malicious peers' bad reputations. Sonja Buchegger and Jean-Yves Le Boudec designed a distributed reputation system using a Bayesian approach, in which the second-hand reputation rating is accepted only when it isn't compatible with the primary rating.³

Several universities are working on research projects involving trust management in P2P applications. At Georgia Tech, Li Xiong and Ling Liu have developed the

PeerTrust model.⁴ Their model is based on a weighted sum of five peer feedback factors: *peer records*, *scope*, *credibility*, *transaction context*, and *community context*. PeerTrust is fully distributed, uses overlay for trust propagation, public-key infrastructure for securing remote scores, and prevents peers from some malicious abuses.

At Stanford University, Hector Garcia-Molina and colleagues proposed the EigenTrust algorithm,⁵ which captures peer reputation in the number of satisfactory transactions and then normalizes it over all participating peers. The algorithm aggregates the scores by a weighted sum of all raw reputation scores. EigenTrust is fully distributed using a DHT-based overlay network. The system also assumes pre-trust peers and uses majority voting to check faulty reputation scores reported.

Weisong Shi and colleagues at Wayne State University proposed the TrustWare system (<http://mist.cs.wayne.edu/trustware.html>),⁶ a trusted middleware for P2P applications. Their approach consists of two models: the Multiple Currency Based Economic model (M-CUBE) and the Personalized Trust model (PET). The M-CUBE model provides a general and flexible substrate to support high-level P2P resource-management services. PET derives peer trustworthiness from long-term reputation evaluation and short-term risk evaluation.

At University of Maryland, Seungjoon Lee and colleagues proposed in the Nice project a scheme for trust inference in P2P networks (www.cs.umd.edu/projects/nice/).⁷

The trust inference consists of two parts for local trust inference and distributed search. After each transaction, the system generates cookies to record direct trust between peers. It also uses trust graphs to infer transitive trust along a peer chain.

References

1. P. Resnick and R. Zeckhauser, "Trust among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System," *The Economics of the Internet and E-commerce*, M.R. Baye, ed., Elsevier, 2002, pp. 127–157.
2. R. Guha et al, "Propagation of Trust and Distrust," *Proc. World Wide Web Conf. (WWW 2004)*, ACM Press, 2004, pp. 403–412.
3. S. Buchegger and J.-Y. Le Boudec, "A Robust Reputation System for P2P and Mobile Ad-Hoc Networks," *Proc. 2nd Workshop Economics of Peer-to-Peer Systems*, 2004; <http://citeseer.ist.psu.edu/buchegger04robust.html>.
4. L. Xiong and L. Liu, "PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities," *IEEE Trans. Knowledge and Data Eng.*, vol. 16, no. 7, 2004, pp. 843–857.
5. S. Kamvar, M. Schlosser, and H. Garcia-Molina, "The EigenTrust Algorithm for Reputation Management in P2P Networks," *Proc. World Wide Web Conf. (WWW 2003)*, ACM Press, 2003, pp. 640–651.
6. Z. Liang, and W. Shi, "PET: A Personalized Trust Model with Reputation and Risk Evaluation for P2P Resource Sharing," *Proc. 38th Ann. Hawaii Int'l Conf. System Sciences*, IEEE CS Press, 2005, pp. 201.2.
7. S. Lee, R. Sherwood, and B. Bhattacharjee, "Cooperative Peer Groups in NICE," *Proc. IEEE Conf. Computer Comm. (INFOCOM 03)*, IEEE CS Press, 2003, pp. 1272–1282.

ty, fuzziness, and incomplete information adaptively. To this end, we first analyzed the eBay auction-based transaction trace data to sort out client behavioral characteristics. Then we built FuzzyTrust, a prototype P2P reputation system that helps establish mutual trust among strangers in P2P transaction applications.

Our system uses fuzzy logic inference rules to calculate local trust scores and to aggregate global reputation. This system benefits from the distinct advantages of fuzzy inferences, which can handle imprecise linguistic terms effectively. In the

"Fuzzy Logic Inference and Applications" sidebar, we give details on how to perform fuzzy inference in the FuzzyTrust system. The system uses distributed-hash-table (DHT) overlay networks to perform fast and secure reputation dissemination among peers.⁴

We evaluated the FuzzyTrust system via simulation experiments over the eBay trace data, which we collected from the public domain during March 2005. We then compared FuzzyTrust with the EigenTrust algorithm proposed at Stanford University,⁵ which calculates global peer reputation by

Fuzzy Logic Inference and Applications

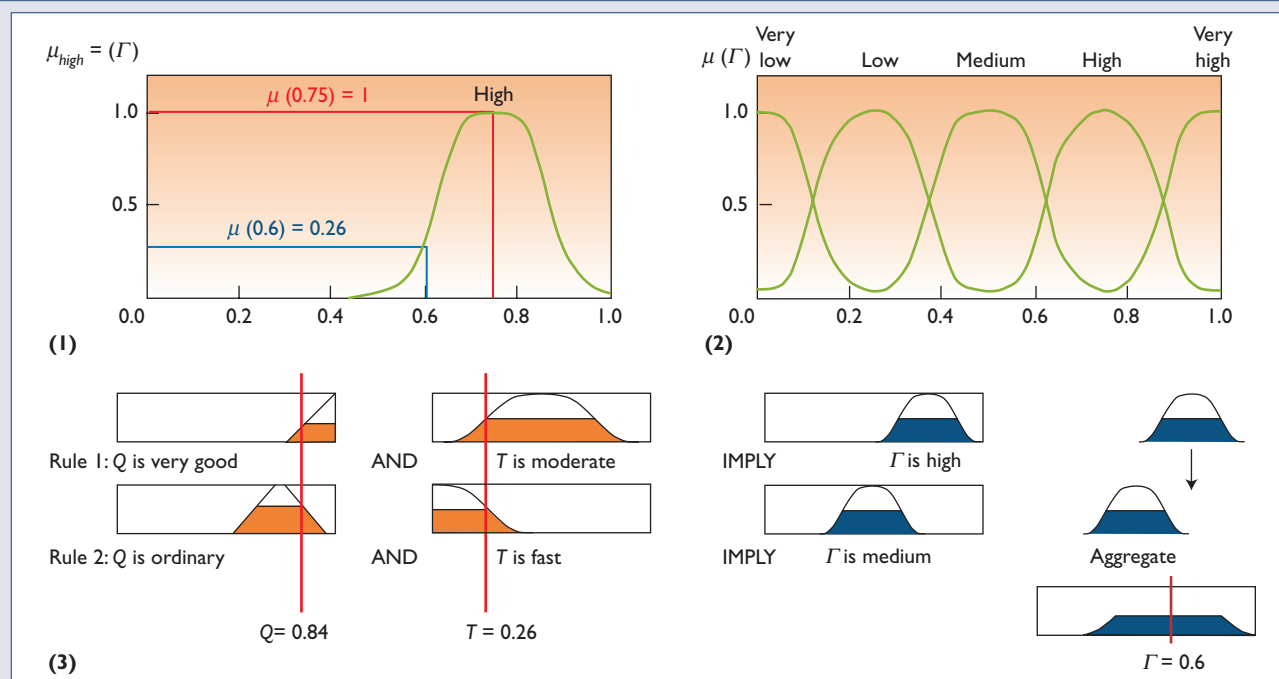


Figure A. Fuzzy membership functions and the fuzzy reputation aggregation procedure. We show (1) the high membership function of a local score (Γ), (2) the five levels of membership functions of Γ , and (3) the application of two rules to induce the seller's local score.

Ever since Lotfi Zadeh proposed fuzzy logic 40 years ago,¹ fuzzy theory has demonstrated its power in managing

uncertainties and mimicking the human decision-making process. (You can find archival articles and tutorials on fuzzy

logic at www.austinlinks.com/Fuzzy/, and books, journals, and research groups on the topic at www.abo.fi/~rfuller/fuzs.html).

considering the entire system's history. We demonstrate via experimentation that FuzzyTrust is more effective in identifying malicious peers and is efficient with reduced message overhead in the global reputation aggregation process.

P2P Reputation Systems for e-Transactions

P2P systems are client-oriented, and the peers involved are autonomous. Such systems can grow or shrink dynamically with self-organizing capabilities. Unfortunately, dynamic features of P2P systems make them prone to cyberspace abuses and attacks. P2P reputation systems aim to stop such abuses and avoid unnecessary losses for well-behaved peers.¹ Figure 1 shows the key components of a typical P2P reputation system.

At the top of the diagram, we identify four major P2P applications: consumer product exchange, distributed file-sharing, security infor-

mation sharing, and processing-cycle sharing among distributed clients. To support consumer product exchange and share security information, we designed our reputation system specifically for e-commerce applications. In such applications, both sellers and buyers are participating peers: the sellers post their merchandise on the Internet for sale, and the buyers place their orders from the client hosts at the Internet edge. We built a P2P overlay network on top of the peer hosts that acts as a virtual network for peer reputation evaluation and dissemination. We preferred using a structured P2P overlay such as DHT rather than unstructured P2P systems because it shortens the data-search process and reduces routing complexity.⁴ In Figure 1, solid lines link the intended exchange partners and symbolize the establishment of mutual trust. The dashed lines refer to intermediate peers involved in the reputation-aggregation process.

Fuzzy Logic Inference and Applications (cont.)

Researchers report successful application of fuzzy theory for adaptive control in robotics, tracking, and consumer electronics; information retrieval in database management; pattern recognition in advanced automation; and fuzzy decision support to handle uncertainties in large-scale information systems.²

We introduce yet another important application to the network security area: fuzzy reputation management. This is a new approach to supporting approximated reasoning. The FuzzyTrust model is useful in manipulating imprecise or uncertain information. Indeed, there are five common features between a reputation system and a fuzzy controller designed for control optimization:

- *Imprecise inputs:* Transaction records can contain untruthful information — a situation similar to the handling of noisy signals in a control application.
- *Linguistics expertise knowledge:* In evaluating reputation, human knowledge must be the fundamental input. This is similar to the use of linguistics information in a control application.
- *Prior information:* We need to use

imprecise transaction records accumulated in the past. This resembles the learning process in evolving the fuzzy controller rule set.

- *Dynamic environment:* Peer behavior changes dynamically, and such changes must be tracked efficiently. This parallels the handling of dynamically changing environmental signals in a control application.
- *Capture of feedback:* We must use feedback from previous reputation evaluation results in a way similar to calibrating a fuzzy controller by system feedback.

To explain basic fuzzy concepts, we use the seller's local score inference example. In fuzzy theory, the *membership function* $\mu(x)$ for a fuzzy variable x specifies the degree of an element belonging to a fuzzy set. It maps x into the range $[0, 1]$, where 1 is full membership and 0 is no membership. Figure A1 shows a high membership function for modeling the local score (Γ), and Figure A2 shows the five levels of membership function. Figure A3 illustrates the inference process. Consider two fuzzy variables: one is the *product quality* (Q) and another is

the *delivery time* (T), with initial values $Q = 0.84$ and $T = 0.26$.

To illustrate, we apply the following two simple fuzzy inference rules in Figure A3:

1. If Q is very good AND T is moderate, then Γ is high.
2. If Q is ordinary AND T is fast, then Γ is medium.

We infer all rules in parallel and determine the resulting membership by assessing all terms in the premise. We apply the fuzzy operator AND to determine the support degree of the rules, and the AGGREGATE operator superimposes two resulting membership curves. We generate the final local score $\Gamma = 0.6$ by defuzzifying from the aggregation result, taking the centroid of the superimposed membership curve in Figure A3. In a real-life P2P reputation system, this fuzzy logic inference process could demand anywhere from tens to hundreds of rules.

References

1. L. Zadeh, "Fuzzy Sets," *Information and Control*, vol. 8, 1965, pp. 338–353.
2. G.J. Klir, U. St. Clair, and B. Yuan, *Fuzzy Set Theory: Foundations and Applications*, Prentice Hall, 1997.

P2P Transaction Applications

In this article, we concentrate on P2P transactions that don't use external authorities to mediate transaction flows. Every peer node is both a client and a server and acts as an information source. In consumer product exchanges, buyers like to know a seller's reputation before they commit to paying for desired goods. In a file-sharing application, downloading peers prefer to receive files from trusted peers to ensure that such files aren't infected with viruses or embedded with Trojan horses. Consumer product exchanges and P2P file sharing usually occur at client desktops.

Conversely, benefiting from recent advancements in Web services architectures, some P2P activities occur at the server level, such as security information sharing. Clearly, a server would always want to receive and process security updates from other trusted servers. In a P2P grid environment, users share their processing cycles to

perform large-scale, collective, and cooperative computational applications. For security reasons, users would want to avoid executing their jobs on remote machines, especially if they couldn't fully trust those machines.⁶

Reputation Evaluation and Dissemination

Reputation is hard to quantify because many dynamic factors are involved. In a fully distributed P2P system involving numerous peers, a peer often can't assess another peer's reputation effectively, but rather must rely on collective opinions from other peers. However, this introduces new challenges, such as how to determine the accuracy of the collected opinions and how to aggregate the conflicting opinions to yield a *global reputation*. Furthermore, the global reputation aggregation process can be very time consuming and incur a heavy messaging overhead when high accuracy or updated reputation are desired.

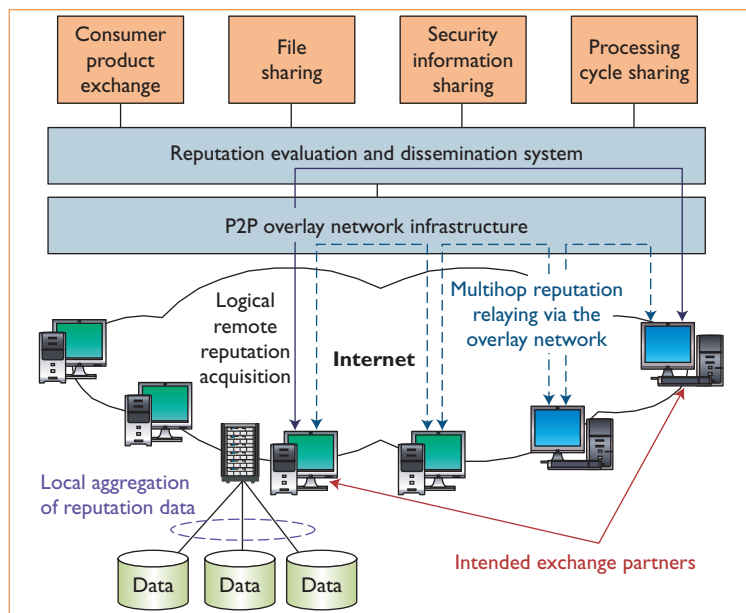


Figure 1. A typical peer reputation aggregation system. The system is built with a P2P overlay network to facilitate trusted commodity exchanges over the Internet. The lower solid line links the intended exchange partners. The upper solid line shows the establishment of mutual trust between the peers. The dashed lines refer to intermediate peers involved in the aggregation process.

An effective P2P reputation system assesses reputation locally and aggregates the meta-reputation globally. It must accurately capture and track various local parameters – buyers should be able to capture potential sellers’ credit records, for example – or preferences with respect to the goods being sold. The reputation system must also accurately filter out untrustworthy second opinions from any malicious peers trying to blacken the reputations of some well-behaved peers in the system.

Analysis of the eBay Transaction Data

eBay (www.ebay.com) is by far the most successful cyberexchange platform currently on the Internet, and most transactions via the site are conducted among strangers. Although eBay has a centralized reputation system, all transactions among the peers are distributed at the client end. The eBay transaction records provide a real-life data set for testing a P2P reputation system’s efficacy.

To obtain eBay transaction records, we extracted public-domain data released on its Web site, crawling over more than 10,000 Web pages from users to produce the trace file. We observed three important characteristics of eBay transac-

tions: super versus small users; unstable transaction intervals for small users; and skewed transaction amounts. We briefly describe these characteristics, and Figure 2 illustrates the peer feedback distribution, user interaction pattern, transaction interval, and variation of transaction amounts illustrated.

Super Users versus Small Users

Figure 2a demonstrates how we sort users in a decreasing order based on how many transactions they’ve conducted. The eBay transactions follow a power-law distribution across the peer ranks. In our analysis, we found a hot-spot situation in which a few super users (big circles in Figure 2b) contributed most of the transactions, whereas many small users (small circles in Figure 2b) engaged in only a few transactions. The super users have a high node degree – or many links with other users, as Figure 2b shows – and the small users have a lower node degree, indicating that they have just a few transaction partners.

Unstable Transactions by Small Users

The time interval between two adjacent transactions is short and stable for super users. For small users, however, the transaction intervals are long and unstable, as Figure 2c illustrates. Our eBay trace shows that some super users conducted more than 10,000 transactions in three months, whereas the small users conducted only a few tens of transactions in several years. Additionally, super users conducted 65 percent of subsequent transactions within 10 minutes of the previous transaction, whereas small users conducted 50 percent of subsequent transactions 40 hours later.

Skewed Transaction Amount

The transaction amount is highly skewed in the sense that a small number of large transactions coexist with a large number of small transactions. The large transactions dominate the total amount transacted. Figure 2d shows typical transaction amounts in the eBay system. Here, the *normalized transaction rank* indicates the percentage of top-valued commodities transacted. We plot the *percentage of cumulative amount* versus the transaction rank. For example, Figure 2d shows that 30 percent of the large transactions make up 70 percent of the total commodity value. This implies that 70 percent of the remaining smaller transactions contribute only 30 percent of the total transaction amount.

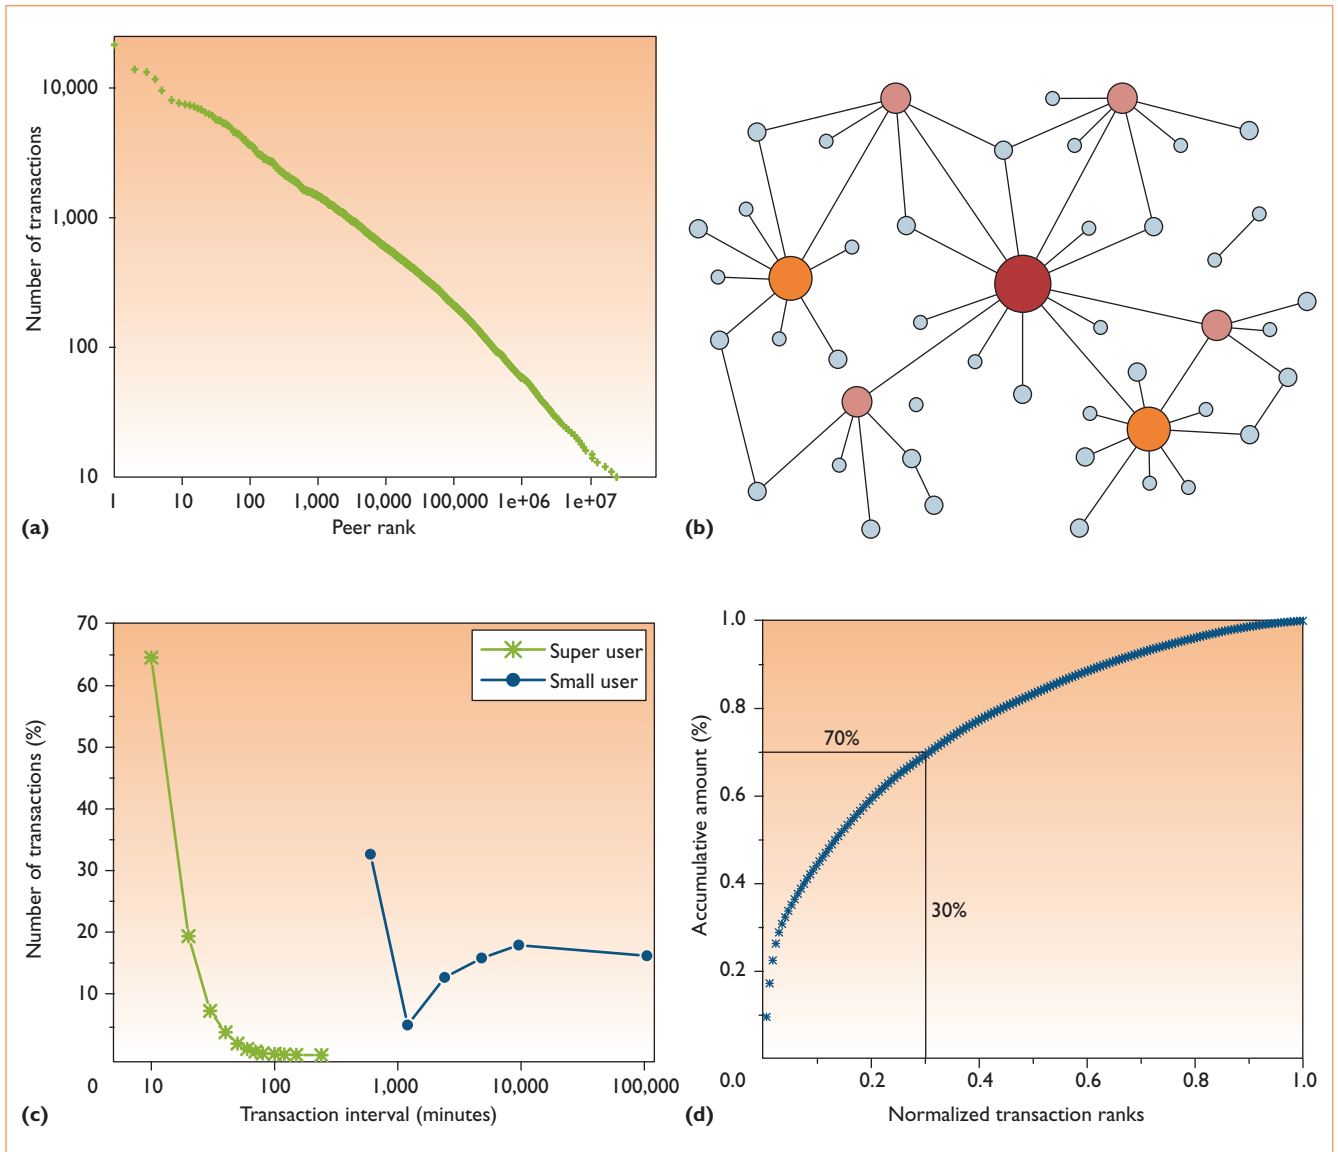


Figure 2. P2P interaction characteristics. We revealed four characteristics using real-life eBay transaction traces. We looked at (a) peer ranks (power-law distribution), (b) hot spots (super versus small users), (c) the transaction interval, and (d) the transaction amount distribution.

FuzzyTrust System Architecture

After analyzing the characteristics of eBay’s transaction data, we developed the FuzzyTrust prototype system for evaluating peer reputation in P2P transactions. We built FuzzyTrust with the fuzzy-logic inference technique introduced in the sidebar. In particular, the system is capable of handling imprecise or uncertain information collected from the peers.

System Design Requirements

Using what we learned from eBay’s transaction characteristics, we suggest three important design criteria:

- First, the network bandwidth consumption required to exchange local trust scores for hot spots can be extremely high. Thus, a reputation system for e-transactions should consider the unbalanced transactions among users.
- Second, to address the lesser impact from small users, a reputation system shouldn’t apply the same evaluation cycle for all peers. The super users should be updated more often than the small users.
- Third, with a skewed transaction amount, it makes sense to evaluate the large transactions more often than the small ones.

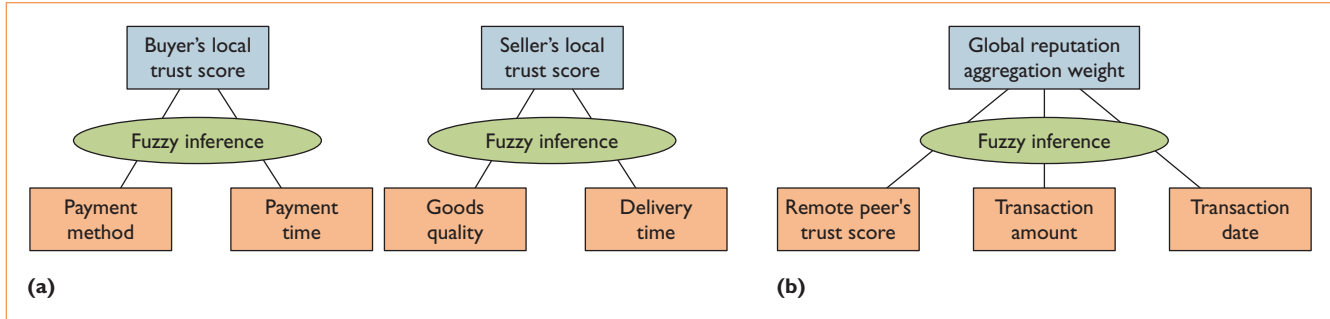


Figure 3. Local trust inference versus global reputation aggregation. The FuzzyTrust system (a) performs fuzzy logic inferences to determine the local trust scores and (b) uses accumulated local scores for weight inference in global reputation aggregation.

Our system works by performing two major inference steps: *local-score calculation* and *global reputation aggregation* (see Figure 3).

Local-Score Computation. In FuzzyTrust, peers perform fuzzy inference on local parameters to generate the local scores. Figure 3a illustrates the local-score calculation for eBay transactions. The fuzzy inference mechanism can capture some uncertainties and is self-adjusting. It can adaptively track the variation of local parameters, such as payment method and time, goods quality, delivery time, and so on. Additional examples of this process are available elsewhere.⁷

Global Reputation Aggregation. The FuzzyTrust system aggregates local trust scores collected from all peers to produce a global reputation for each peer. The system uses fuzzy inference to obtain the global reputation aggregation weights, as illustrated in Figure 3b. We determine the aggregation weights using three variables: the peer's reputation, the transaction date, and the transaction amount. In a full-scale P2P reputation system, the number of fuzzy inference rules could be extended to several hundreds. However, we list five frequently used fuzzy inference rules we applied to the prototype FuzzyTrust system construction.

1. If the transaction amount is very high and the transaction time is new, then the aggregation weight is very large.
2. If the transaction amount is very low or the transaction time is very old, then the aggregation weight is small.
3. If a peer's reputation is good and the transaction amount is high, then the aggregation weight is very large.

4. If a peer's reputation is good and the transaction amount is low, then the aggregation weight is medium.
5. If a peer's reputation is bad, then the aggregation weight is very small.

We calculate the global reputation using the following formula:

$$R_i = \sum_{j \in S} \left(\frac{w_j}{\sum_{j \in S} w_j} t_{ji} \right) = \frac{\sum_{j \in S} w_j t_{ji}}{\sum_{j \in S} w_j}, \quad (1)$$

where R_i is the global reputation of peer i , S is the set of peers with whom peer i has conducted transactions, t_{ji} is the local trust score of peer i rated by peer j , and w_j is the aggregation weight of t_{ji} . The global aggregation process runs multiple iterations until each R_i converges to a stable global reputation rating for peer i .

DHT-Based Overlay Implementation

We implemented the prototype FuzzyTrust system on a DHT-based P2P overlay network, with an architecture similar to that of Chord,⁴ a DHT ring that provides fast trust aggregation and secure message transmission. The Chord system is highly scalable, robust to failure, and self-organizing in that it handles peer join and leave from the system. Figure 4 shows the DHT-based FuzzyTrust system architecture.

Each peer maintains two tables: a *transaction record table* to maintain transaction records with remote peers, and a *local score table* to maintain remote peers' evaluated trust scores. Based on the transaction records, we infer the global aggregation weights through the fuzzy inference system. When performing global rep-

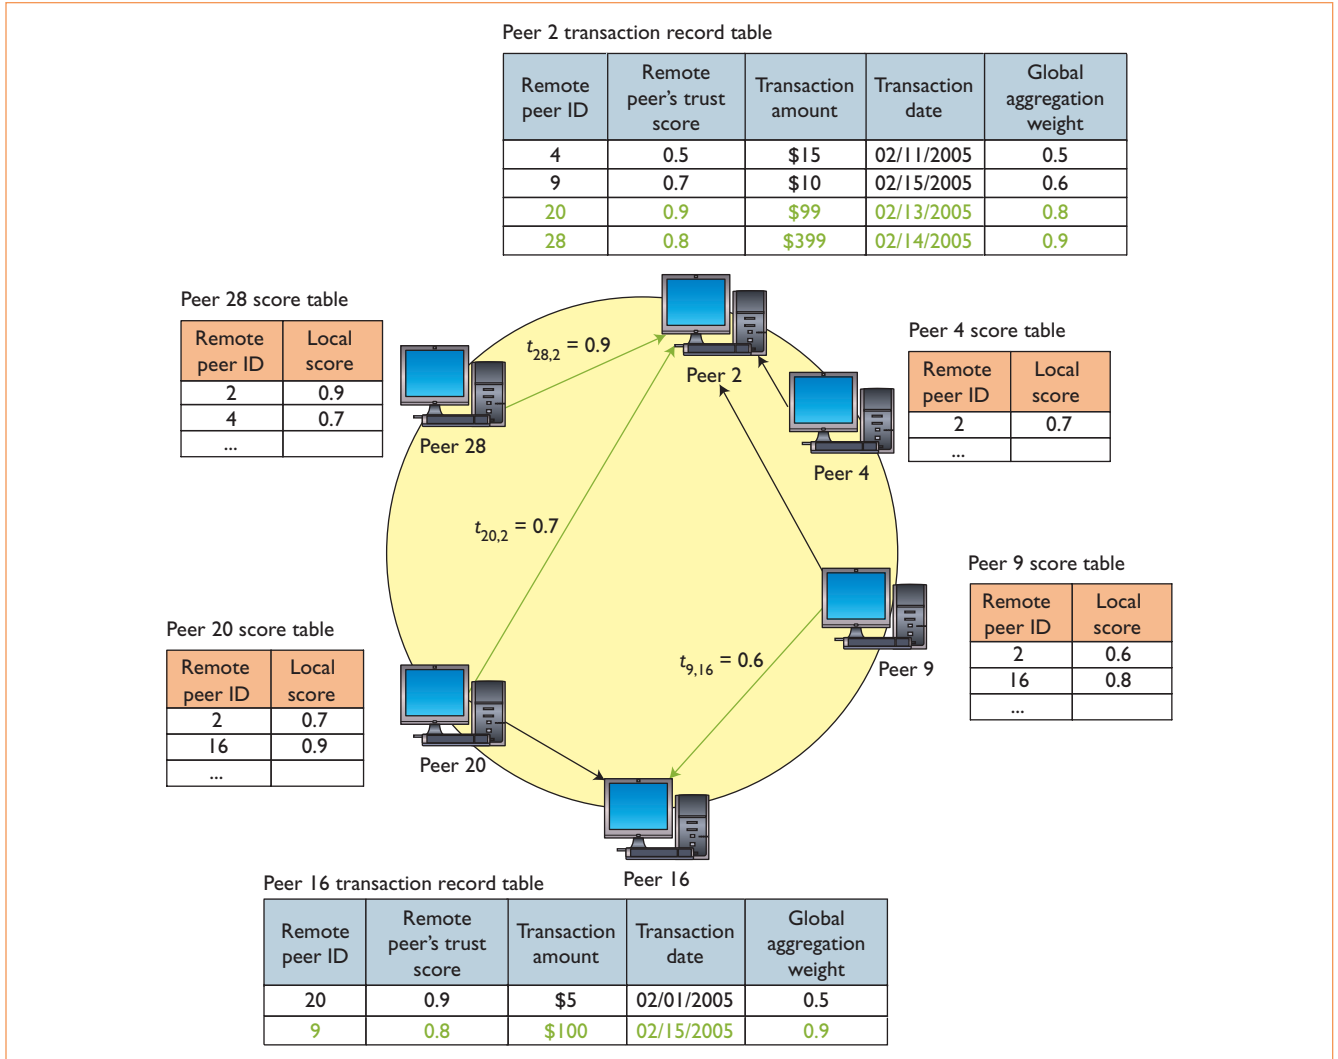


Figure 4. The FuzzyTrust system architecture. We built FuzzyTrust with a DHT-based overlay network supporting the local trust score dissemination and global reputation aggregation.

putation aggregation, each peer queries the trust scores from remote peers. To tackle the hot-spot issue, the system partially queries qualified peers that meet an aggregation threshold. Figure 5 shows an example of global reputation aggregation based on the DHT configuration in Figure 4.

In this particular example, the system sets an aggregation threshold of 0.7 for Peer 2. Thus, Peers 28 and 20 are qualified to be queried, whereas Peers 9 and 4 aren't qualified and thus aren't queried. Peers 28 and 20 transmit their local scores $t_{28,2}$ and $t_{20,2}$ to Peer 2. We use Equation 1 to calculate Peer 2's global reputation. We set a high threshold for the hot spots, so the system queries only a few peers, thus reducing the bandwidth consumption.

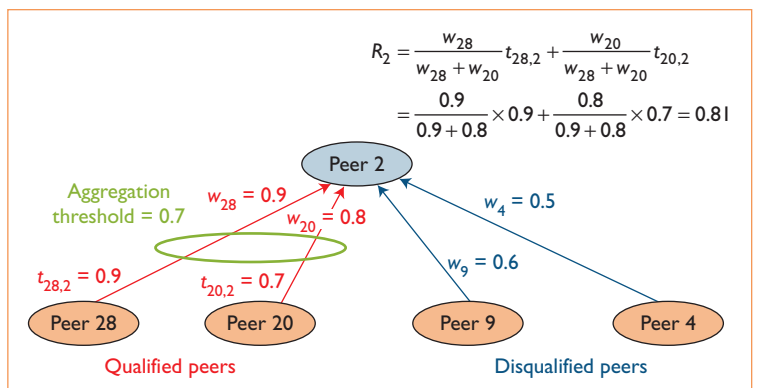


Figure 5. Global reputation aggregation using the DHT ring from Figure 4. The system sets Peer 2 with an aggregation threshold of 0.7. Thus Peer 9 and Peer 4 are disqualified from participating in the evaluation process. Peer 2's reputation is aggregated from two local trust scores, $t_{28,2}$ and $t_{20,2}$, generated by Peer 28 and Peer 20, respectively.

Simulated Performance Results

We evaluated the FuzzyTrust and EigenTrust systems for P2P transactions using the eBay trace data set we collected at USC. We compare the two P2P reputation systems in the following three performance matrices:

- the *convergence time* needed to establish each peer's global reputation,
- the *detection rate* of malicious peers, and
- the *message overhead* involved in global reputation aggregation.

We simulated the global reputation aggregation of the two P2P reputation systems over P2P systems with $N = 100$ to 10,000 peers. For each case, the super users or most active peers performed $10N$ transactions. The number of transactions per peer follows the power-law distribution with a slope of -1.215 , which we obtained from the eBay trace experiments. The transaction date follows a Poisson distribution with an arrival rate $\lambda = 0.2$ transactions per minute. The simulation calculates all local trust scores and aggregates global peer reputation. Figure 6 plots the simulated performance results in four aspects: global convergence time, malicious-peer-detection rate, messaging overhead for individual peers, and total messaging overhead.

Global Convergence Time

We plot the convergence times of the EigenTrust and FuzzyTrust systems in Figure 6a. Their convergence times are quite close to each other, although the FuzzyTrust system converges slightly faster than the EigenTrust system. Overall, their convergence times are comparable, both increasing linearly with respect to the P2P network size.

Malicious Peer Detection Rate

Malicious peers are those that, as buyers, make late or no payments, and, as sellers, deliver bad-quality goods or fail to deliver goods at all. Let α be the percentage of malicious peers in the P2P system and m be the number of malicious peers in the system. We have $m = \alpha N$, where N is the system size. In our simulation experiments, we preset $\alpha = 0.3$, meaning 30 percent of the peers are potentially abusing the system. We denote the number of malicious peers detected as $d(t)$, which depends on the time of measurement. We define the *detection rate* θ of malicious peers as

$$\theta(t) = d(t) / m = d(t) / \alpha N. \quad (2)$$

Figure 6b plots the measured malicious-peer-detection rate as a function of time represented by the number of iterations in the evaluation process. The bars correspond to various system sizes N with a fixed $\alpha = 0.3$. We observed two important phenomena: for all system sizes, FuzzyTrust detects more than 99 percent of malicious peers after four iterations, and malicious peers are much easier to detect in small systems. Given a system with 100 peers, for example, the FuzzyTrust system detects more than 90 percent of malicious peers after the first iteration. The EigenTrust system has slightly lower detection rates, and isn't plotted in Figure 6b.

Messaging Overhead

Figure 6c plots the *average message overhead* per individual peer for a system of 1,000 peers. We rank peers by number of messages transmitted. Compared with the EigenTrust system, messages transmitted in FuzzyTrust are more evenly distributed among the peers. On average, each peer in FuzzyTrust transmits 17 messages with a standard deviation of 15, compared with 73 messages transmitted in EigenTrust with a standard deviation of 227.

Using the FuzzyTrust system, the peers do experience a noticeably lower messaging overhead. In other words, our system can better alleviate the hot-spot problem, whereas the EigenTrust system cannot. Figure 6d shows the total number of messages required to reach a global convergence. The gap between total message overheads widens as the network size increases. The FuzzyTrust system is scalable in handling an even larger number of P2P services. For 10,000 peers, our system needs 180,000 messages to aggregate all global reputation scores, whereas the EigenTrust system need 580,000 messages to perform the same task.

Fuzzy logic inference is clearly effective for distributed trust management in P2P networks. Based on the eBay trace data, our experimental results show that both the FuzzyTrust and EigenTrust systems have comparable global reputation convergence time. Our system detects all malicious peers after a few aggregation iterations. On average, the peers transmit fewer messages using the FuzzyTrust system than using EigenTrust when performing global reputation aggregation.

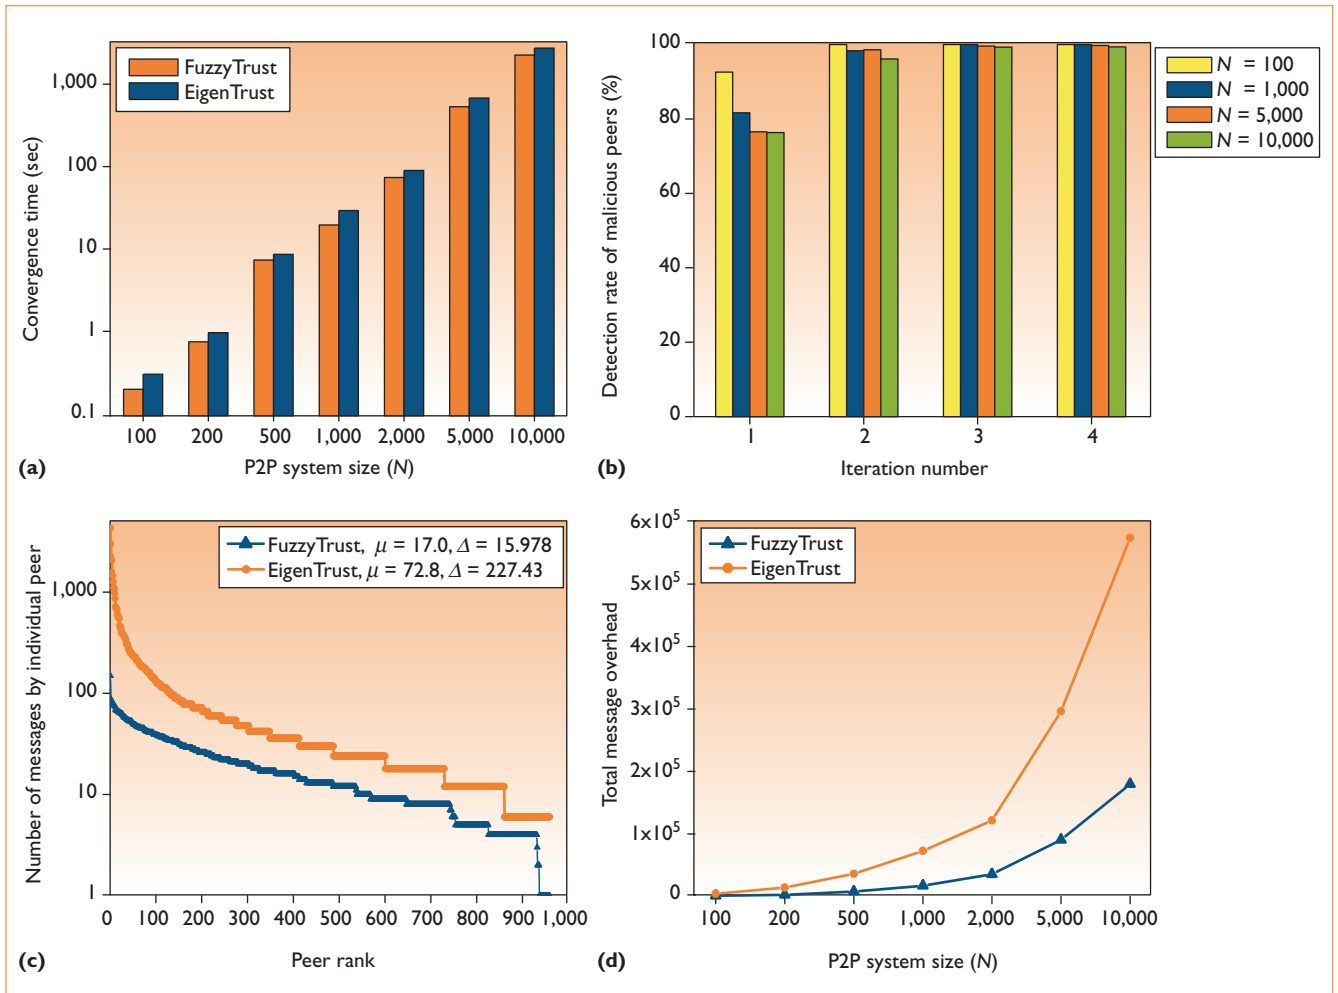


Figure 6. Simulated performance of two P2P reputation systems. We compared the FuzzyTrust system with the EigenTrust system over the same eBay transaction data set. We plot (a) the global convergence time, (b) the detection rate of malicious peers, (c) the message overhead per individual peer, and (d) the total message overhead needed for global reputation aggregation.

Through our experiments, we reveal the power-law connectivity of the DHT-based overlay network for fast establishment of global trust among a scalable number of peers.

For the future, we suggest extending the work on P2P reputation systems by considering peer anonymity and secure storage of global reputations.⁹ Evaluating various P2P systems over distributed file sharing and P2P grid computing is another direction for future research.⁶ For the FuzzyTrust system, we could consider additional e-commerce data traces other than eBay’s, such as that from Amazon or Yahoo transaction records. We might also further evaluate system performance via more real-life experiments. It would be useful to see the FuzzyTrust prototype converted as a production system in real-life P2P applications.¹⁰ □

Acknowledgments

US National Science Foundation Information Technology Research grant NSF 0325409 supported this work. We thank the USC GridSec research team for their technical assistance and useful discussions.

References

1. S. Marti and H. Garcia-Molina, “Taxonomy of Trust: Categorizing P2P Reputation Systems,” *Computer Networks*, special issue on trust and reputation in P2P systems, 2005.
2. D.W. Manchala, “E-Commerce Trust Metrics and Models,” *IEEE Internet Computing*, vol. 4, no. 2, 2000, pp. 36–44.
3. L. Xiong and L. Liu, “PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities,” *IEEE Trans. Knowledge and Data Eng.*, vol. 16, no. 7, 2004, pp. 843–857.

4. I. Stoica et al., "Chord: A Scalable Peer-to-Peer Lookup Protocol for Internet Applications," *IEEE/ACM Trans. Networking*, vol. 11, no. 1, 2003, pp. 17–32.
5. S. Kamvar, M. Schlosser, and H. Garcia-Molina, "The EigenTrust Algorithm for Reputation Management in P2P Networks," *Proc. World Wide Web Conf. (WWW 2003)*, ACM Press, 2003, pp. 640–651.
6. R. Zhou and K. Hwang, *Trust Overlay Networks for Global Reputation Aggregation in P2P Grid Computing*, tech. report TR-2005-16, Internet and Grid Computing Lab., Univ. of Southern California, October 2005.
7. S. Song, K. Hwang, and Y.-K. Kwok, "Trusted Grid Computing with Security Binding and Trust Integration," *J. Grid Computing*, vol. 3, no. 1, 2005; www.springerlink.com/link.asp?id=7028153w8653340.
8. S. Marti and H. Garcia-Molina, "Limited Reputation Sharing in P2P Systems," *Proc. ACM Conf. Electronic Commerce (EC 04)*, ACM Press, 2004, pp. 91–101.
9. A. Singh and L. Liu, "TrustMe: Anonymous Management of Trust Relationships in Decentralized P2P Systems," *Proc. 3rd Int'l Conf. Peer-to-Peer Computing (P2P 2003)*, IEEE CS Press, 2003, pp. 142–149.
10. S. Song, "Trusted Grid and P2P Computing with Security Binding and Reputation Aggregation," PhD dissertation, Dept. Computer Science, Univ. of Southern California, Nov. 2005.

Shanshan Song is a PhD student in computer science at the University of Southern California. She has worked as a research assistant at the USC Internet and Grid Computing Laboratory. Her technical expertise lies in grid computing,

network security, and trust management in P2P and grid systems. Song has a BS in computer science from the University of Science and Technology of China. She is a student member of the IEEE and the ACM. Contact her at shanshan.song@usc.edu.

Kai Hwang is a professor of electrical engineering and computer science at the University of Southern California, director of the USC Internet and Grid Computing Laboratory, and leader of USC's GridSec project. He specializes in computer architecture, parallel processing, network security, distributed computing, and grid and P2P systems. Hwang has a PhD in electrical engineering and computer science from the University of California, Berkeley. He is a fellow of the IEEE. Contact him at kaihwang@usc.edu; <http://GridSec.usc.edu/Hwang.html>.

Runfang Zhou is a PhD student in computer science at the University of Southern California. Her research interests include P2P reputation systems, overlay networks, and network security. Zhou has a BS and an MS in computer science from Southeast University, China. Contact her at rzhou@usc.edu.

Yu-Kwong Kwok is an associate professor of electrical and electronic engineering at the University of Hong Kong. His research interests include grid and mobile computing, wireless communications, and network protocols. Kwok has a PhD in computer science from the Hong Kong University of Science and Technology. He is a senior member of the IEEE. Contact him at ykwok@hku.hk.

IEEE Transactions on Mobile Computing

A revolutionary new quarterly journal that seeks out and delivers the very best peer-reviewed research results on mobility of users, systems, data, computing information organization and access, services, management, and applications. *IEEE Transactions on Mobile Computing* gives you remarkable breadth and depth of coverage ...



Architectures
Support Services
Algorithm/Protocol Design and Analysis
Mobile Environment
Mobile Communication Systems
Applications
Emerging Technologies



To subscribe:
<http://computer.org/tmc>
or call
USA and CANADA:
+1 800 678 4333
WORLDWIDE:
+1 732 981 0060