Deposit and Insurance Schemes for Non-Monitored P2P Service

Rakesh Kumar
Department of Electrical and Computer Engineering,
Polytechnic University, Brooklyn, NY
rkumar04@utopia.poly.edu

Keith W. Ross
Department of Computer and Information Science,
Polytechnic University, Brooklyn, NY
ross@poly.edu

***Abstract*— We study mediated P2P marketplaces in which payment for P2P service is done via a trusted intermediary but the intermediary cannot monitor the rendered service. The lack of direct monitoring provides opportunities for malicious peers to launch crippling "lying attacks" against the marketplace. In this paper we propose a suite of payment schemes that are specifically designed to thwart lying attacks. In particular, we explore the use of deposits and insurance in our payment schemes, and determine bounds on satisfactory deposit and insurance levels. The goal of this paper is to introduce to the research community the problem of designing robust payment schemes for mediated P2P marketplaces and to stimulate further research in the area.**

I. Introduction

In this paper we study P2P marketplaces with three physical components: service provider peers; service consumer peers; and a trusted mediator. In these P2P marketplaces, the provider peers render service directly to the consumer peers. The mediator acts as a trusted payment intermediary between provider peers and consumer peers - that is, although the provider peers directly provide service to the consumer peers, the payment for service passes through the mediator. We refer to such P2P marketplaces as **mediated P2P marketplaces**.

A salient characteristic of a mediated P2P marketplace is that the services rendered are non-monitored – it is impossible for the mediator to unambiguously ascribe blame when either the provider or the consumer claims that there has been a service fault. In particular, in a service transaction between a provider peer and a consumer peer, the provider peer may claim that it rendered service while the consumer peer may claim it didn't receive any service.

A. Peer-Powered Content Distribution

As a concrete example of a mediated P2P marketplace, consider peer-powered content distribution, which is a P2P version of a Content Distribution Network (CDN). Recall that in the CDN paradigm, there are three parties: (i) the clients who want to consume content; (ii) the content owners (e.g., CNN, the BBC, Microsoft), who want to sell their content to clients; and (iii) the CDN company (e.g., Akamai), which distribute the the owners' content to the clients. A CDN company typically provides its content distribution service by installing hundreds or even thousands of servers throughout the Internet, replicating the owners' content on its servers, redirecting client requests to one of its servers, and sending the requesting content directly from the CDN server to the client. One drawback of this paradigm is that it requires the installation and maintenance of an infrastructure of hundreds of CDN servers.

A **peer-powered content distribution**, which provides a (potentially) lower-cost alternative to a traditional CDN. The basic idea is to replace the CDN servers with user peers. Thus, in peer-powered content distribution, the owners' content is stored in user peers rather than in dedicated servers, and the user peers send the content directly to the clients [1]. The user peers are personal computers with Ethernet or residential broadband access to the Internet. We refer to the user peers that send the content as "distributors" and the clients that receive the content as "consumers". The content distribution flow for peer-powered content distribution is shown in Figure 1. The content could be simply an elastic file transfer (e.g., over TCP) or streamed to the consumer, depending on the agreement specified in the contracts.

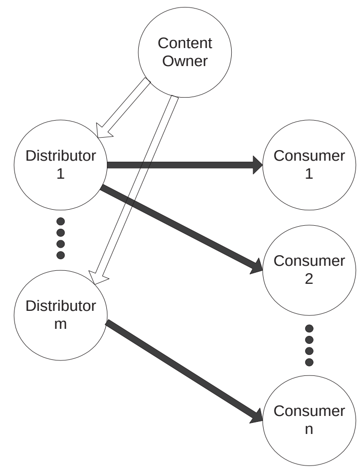

Fig. 1. *Content Flows*

Peer-powered content distribution obviates the need for dedicated servers and bandwidth, providing the potential for significant cost savings. However, for this paradigm to be successful, users must be given incentives so that they agree to distribute content on the behalf of the content owners. This would likely be done via a micropayment scheme, in which the consumers pay the content owner and the content owner pays the distributors for distributing the content, as shown in Figure 2. Thus, two contracts are required [2]. One contract would be between the content owner and the distributor, which specifies how the content owner pays the distributor and the distribution obligations of the distributor; and one contract between the consumer and the content owner, which specifies how the consumer pays the content owner and the service levels that the consumer should expect.

There are two broad approaches for peer-powered content distribution. In the first approach, a content owner (such as CNN) directly commissions peers to distribute its content to consumers. In the second approach, there is a mediator which commissions peer distributors and works for multiple content owners. Such a mediator would be similar to a CDN company, such as Akamai, except that it would not deploy its own dedicated servers, but would instead outsource the content distribution to user peers.

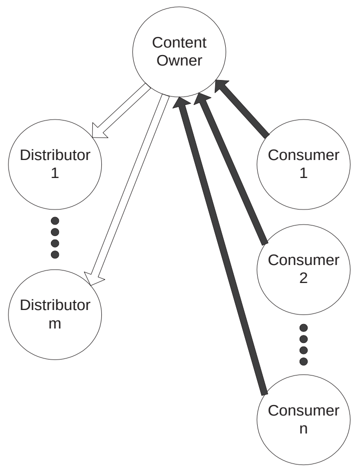

Fig. 2. *Payment Flows*

B. Monitoring service in a mediated P2P Marketplace

Having now described a concrete example, we now return to the context of a mediated P2P marketplace, which more generally consists of a mediator, provider peers and consumer peers. Because of its inherent P2P nature, the mediator cannot directly monitor the service provided by a provider to a consumer. Consequently, the mediator cannot easily resolve any disputes that might occur between a provider and a consumer. In a typical dispute, the consumer may claim that it didn't receive service at the contracted service level and the provider may claim that it rendered the service at the contracted level. The provider can lie about the service it rendered, and the consumer can lie about the service it received.

The inability to monitor is a critical problem in a mediated P2P marketplace. In particular, the lack of direct monitoring provides opportunities for malicious peers to launch crippling "lying attacks" against the marketplace. In this paper we propose a series of payment schemes that are specifically designed to thwart lying attacks in mediated P2P marketplaces. The contributions of this paper are the following:

- We first examine a simple payment scheme in which the consumer peer is required to pay before service is rendered, and the provider peer is paid after service is rendered and only if the consumer peer indicates that it is satisfied with the service. We refer to this scheme has **pre-pay-post-pay**. We argue that this simple payment scheme provides a first-line of defense against basic attacks. However, being susceptible to an attack from conspiring providers, it alone does not prevent large-scale lying attacks.
- To address the insufficiencies of pre-pay-post-pay, we propose using "deposits". Here, both the provider peer and the consumer peer provide a deposit to the mediator before service is rendered. If the transaction is declared a failure, both the provider and consumer forfeit their deposits to the mediator. Such a scheme discourages malicious behavior and provides a defense against lying attack.
- We propose an alternative to deposits in which insurance coverage is available to the provider and consumer peers. In case of transaction failure, the consumer and the provider are refunded the cost of the transaction. The scheme rewards more reliable peers with lower insurance premiums and penalizes malicious peers with higher premiums.

The goal of this paper is to introduce to the research community the problem of designing robust payment schemes for mediated P2P marketplaces and to stimulate further research in the area.

C. eBay

eBay [3] is an example of a marketplace with non-monitored service. Just as in a mediated P2P marketplace, eBay cannot easily resolve disputes between transacting parties. To address this issue, eBay employs a combination of user reputation and user registration of credit card (or bank account) information. eBay successfully employs a user reputation system that aggregates ratings over all of the user's transactions. This reputation system alone is insufficient in combatting the sybil attack [4], since without additional mechanisms, a malicious user could create a fresh identity as soon as it is flagged as malicious. In eBay users are required to register with either credit card or bank account information; they are also required to accept legal agreements about dispute resolution.

In this paper we target light-weight P2P settings where recording financial information to prevent the creation of a large number of identities or using litigation for dispute resolution would be infeasible. In our setting, a peer is permitted to create new identities without limit. Thus the eBay approach is not directly applicable in our setting.

II. Pre-Payment and Post-Payment Schemes

Consider a transaction involving a consumer, a provider and the mediator. The consumer (call her Alice) needs to pay the mediator, and the mediator has to in-turn pay the provider (call him Bob) for providing service to the consumer. The consumer, Alice, can be asked to pay the mediator before or after the service transaction. Similarly, the mediator can pay the provider, Bob, before or after the transaction.

Let us consider a scenario wherein Alice is allowed to post-pay after receiving satisfactory service from Bob. By allowing Alice to post-pay, she has the incentive to falsely claim that she did not receive any or adequate service from Bob, thereby not paying for a service she actually received. Thus, it is clear that post-pay for the consumer is a poor design choice. We therefore require the consumer to pre-pay. This is akin to the buyer paying for the goods in eBay before the goods are delivered. This only leaves the question of what payment strategy to use for the distributor in the system. There are two possible approaches: pre-pay-pre-pay; and pre-pay-post-pay.

Pre-pay-pre-pay is when both payment transactions, between Alice and the mediator and between the mediator and Bob, must occur before any service is provided. Unfortunately, in this scheme Bob has an incentive to lie as he can falsely claim to have delivered the content after collecting payment. Thus, by lying, Bob has a positive payoff of the negotiated price of the transaction.

In pre-pay-post-pay, Alice must pay the mediator before receiving service, and Bob gets paid only after Alice confirms completion of a successful service transaction. It is interesting to draw similarities of this approach with the popular Escrow system [5]. In Escrow, the buyer deposits funds with an escrow service and funds are released to the seller only when the buyer confirms receipt of good service from the seller.

We briefly note that eBay employs a pre-pay-pre-pay strategy. This is because for the use of pre-pay-post-pay, one needs existence of a trusted clearing-house for payments. This is not feasible for eBay because its users use a variety of payment mechanisms, such as personal checks and money orders, and eBay cannot act as a clearing-house for such payments. Also in paypal [6], a payment choice which is quickly growing in popularity in eBay, funds from the buyer are immediately deposited in the seller's account. Thus paypal resorts to a pre-pay-pre-pay approach. In our peer-powered content distribution example, the content owner assumes the responsibility of the clearing-house. Further, the potential weaknesses of eBay's use of pre-pay-pre-pay are offset by its tight control over user identities and by its litigation threat.

A. *Is Pre-Pay-Post-Pay Sufficient?*

To summarize, in pre-pay-post-pay:

- Bob's contractual obligation is to provide service.
- Alice prepays before service is rendered.
- Alice reports if she received the contracted service; if yes, the mediator pays Bob.

We say that Bob is truthful in a service transaction if he renders service; otherwise, Bob is said to lie in the transaction. We say that Alice is truthful in a service transaction if she correctly reports whether she received service from Bob; otherwise, we say that Alice lied in the transaction.

We have argued that pre-pay for the consumer and post-pay for the distributor is better than the other choices available for a mediated P2P marketplace. However, it nevertheless has its flaws

1) *Bob lies*: Bob could claim that he rendered service when he actually did not. Bob does not loose anything by doing this yet succeeds in dissatisfying a consumer peer which has already paid for the service and may therefore never return. This results in losses of future revenue to the mediator. Furthermore, a competitor could create thousands of malicious providers, each with a different name, and each launching this attack.

2) *Alice lies*: She could falsely report that she did not receive service, when she actually did. This particular attack is not as serious. It cannot be used as part of a sibling attack, as the Alices must pay each time they request the service.

Assuming Bob is a rational peer, he will always claim that he provided the service to Alice even if he did not. We say that the mediator declares a transaction to be successful if the Alice claims the transaction to be a success. Otherwise, the transaction is declared a failure.

Within this framework, we distinguish between two different transaction results:

- *actual transaction result:* defined to be a success if Bob provides service and failure otherwise.
- *mediator's declared transaction result:* transaction result declared by the mediator.

It is to be noted that the actual transaction result is only visible to Alice who is not trustable. Therefore she may lie about the actual transaction result and in turn the "mediator's declared transaction result" may not be the same as actual transaction result. Also, note that if neither of the transacting peers lie, the mediator's declared transaction result will be the same as the actual transaction result. In Tables I and II we provide values of both actual and mediator's declared transaction results for the different cases of: (1) Bob is truthful/lying about if he provided service and (2) Alice is truthful/lying about if she received service.

The remainder of the paper is focused on different strategies to solve the deficiencies of the pre-pay-post-pay payment strategy. Specifically, to ensure that Alice and Bob have strong dis-incentives for lying.

	Bob Truthful	**Bob Lies**
Alice Truthful	success	failure
Alice Lies	success	failure

TABLE I
ACTUAL TRANSACTION RESULT

	Bob Truthful	**Bob Lies**
Alice Truthful	success	failure
Alice Lies	failure	success

TABLE II
MEDIATOR'S DECLARED TRANSACTION RESULT

Further, throughout this paper we assume that all failures are caused by the providers and consumers. Network failures, for example, are resolved through multiple service transfer attempts among the consumer and provider.

III. DEPOSITS

Before describing our deposit scheme, we first describe our model of peer behavior in the system. Recent work has used a simple model of peer behavior whereby any given peer acts malicious (lies) with an unknown but constant probability[7],[8], [9]. A Bayesian estimator for this probability is developed in [7]. The estimator uses the record of past success and failures. It also takes into account "reputation fading" in which the most recent results are given more weight.

We believe that accurately estimating the probability that a peer lies in the subsequent transaction strongly depends on the nature of system under consideration. The problem of designing such an estimator is outside the scope of this paper; henceforth, we assume that for every peer in the system, we have available an estimate of the probability of truthful behavior in the next transaction. If the mediator declares the transaction to be a failure then both of the transacting peers are recorded as having lied in the transaction. One would then expect that value of estimators of probability of lying increases.

	Bob Truthful	**Bob Lies**
Alice Truthful	(c, γ_a)	$(c + d_a, 0)$
Alice Lies	$(c + d_a, \gamma_a + \gamma)$	Assumed not possible

TABLE III
(COST,BENEFIT) REPRESENTATION FOR ALICE WITH USE OF DEPOSITS

	Bob Truthful	**Bob Lies**
Alice Truthful	(γ_b, c)	$(0, \gamma - d_b)$
Alice Lies	$(\gamma_b, -d_b)$	Assumed not possible

TABLE IV
(COST,BENEFIT) REPRESENTATION FOR BOB WITH USE OF DEPOSITS

We now propose a simple scheme which provides disincentives for lying. This scheme requires both Alice and Bob to pay refundable deposits of d_a and d_b units, respectively, to the mediator, before service is rendered from Bob to Alice. If the mediator declares the transaction to be a failure because Alice claims that she did not receive service, these deposits are forfeited to the mediator.

The inherent assumption in such a scheme is that Alice is willing to forego her security deposit by truthfully declaring Bob's failure to provide service. Alice's motivation being to retaliate against Bob by making him loose his security deposit. This is because a strictly rational Alice would rather claim a failed transaction as success so that she does not loose her own security deposit d_a even if it means that Bob will not loose his security deposit either. Note that this is not an unrealistic assumption. For instance, in E-Bay transactions, due to "retaliation," rarely do the transacting parties leave each other inconsistent feedback (that is, typically both parties claim success or both parties claim failure). Even then the transacting parties in eBay are willing to leave a negative feedback in case of dissatisfaction with the service. However, we relax this assumption in section IV and present an improved scheme based on insurance.

An interesting question is, how large should Alice's and Bob's deposits be? Clearly, the deposits must be significant enough to discourage lying. At the same time, they should not be so large so that expected cost of playing becomes larger than the expected utility. To calculate appropriate amounts of the deposits, we proceed as follows. Let the negotiated price of the transaction be c units, and Alice's and Bob's estimated probability of truthful behavior in the next transaction be p_a and p_b, respectively. We model the Alice's benefit from watching the video as γ_a and Bob's cost for streaming the video as γ_b units. Let the payoff to Alice or Bob for lying be denoted γ. For Alice we show the cost and benefits involved in different situations in Table III and for Bob in Table IV. We now give a standard cost-benefit analysis to understand how the value of the deposits influence the decision making of both peers.

1) Analysis for Alice: The expected cost to Alice with different behavior strategies are:

$$\begin{aligned} E[\text{Cost to Alice}|\text{Alice truthful}] &= p_b c + (1-p_b)(c+d_a) \\ &= c + d_a(1-p_b) \\ E[\text{Cost to Alice}|\text{Alice lies}] &= p_b(c+d_a) \end{aligned}$$

Similarly, the benefit equations are

$$\begin{aligned} E[\text{Benefit to Alice}|\text{Alice truthful}] &= p_b\gamma_a \\ E[\text{Benefit to Alice}|\text{Alice lies}] &= p_b(\gamma_a+\gamma) \end{aligned}$$

Naturally $\gamma_a p_b \geq c$, for otherwise Alice would not have entered into a contract at a negotiated price of c units at all. Further, we would like to ensure that the deposit value is small enough so that

$$E[\text{Cost to Alice}|\text{Alice truthful}] < E[\text{Benefit to Alice}|\text{Alice truthful}]$$

or equivalently,

$$\begin{aligned} & c + d_a(1-p_b) < p_b\gamma_a \\ & \Rightarrow d_a < \frac{p_b\gamma_a - c}{1-p_b} \end{aligned} \quad (1)$$

On the other hand we would also like to make sure that deposit is large enough so that

$$E[\text{Cost to Alice}|\text{Alice lies}] > E[\text{Benefit to Alice}|\text{Alice lies}]$$

or equivalently,

$$\begin{aligned} & p_b(c+d_a) > p_b(\gamma_a+\gamma) \\ & \Rightarrow d_a > \gamma_a + \gamma - c \end{aligned} \quad (2)$$

From equations (1) and (2)

$$\gamma_a + \gamma - c < d_a < \frac{p_b\gamma_a - c}{1-p_b} \quad (3)$$

2) Analysis for Bob: Bob's cost for different behavior strategy can be written as:

$$\begin{aligned} E[\text{Cost to Bob}|\text{Bob truthful}] &= p_a\gamma_b + (1-p_a)\gamma_b = \gamma_b \\ E[\text{Cost to Bob}|\text{Bob lying}] &= 0 \end{aligned}$$

similarly,

$$\begin{aligned} E[\text{Benefit to Bob}|\text{Bob truthful}] &= p_a c + (1-p_a)(-d_b) \\ E[\text{Benefit to Bob}|\text{Bob lying}] &= \gamma - d_b \end{aligned}$$

Naturally $\gamma_b \leq p_a c$ otherwise Bob would not agree to enter a contract with the negotiated price of c units. Now to ensure that that the deposit value is small enough so that

$$E[\text{Cost to Bob}|\text{Bob truthful}] < E[\text{Benefit to Bob}|\text{Bob truthful}]$$

or equivalently,

$$\begin{aligned} & \gamma_b < p_a c + (1-p_a)(-d_b) \\ & \Rightarrow d_b < \frac{p_a c - \gamma_b}{1-p_a} \end{aligned} \quad (4)$$

To ensure that deposit is large enough for lying Bobs so that

$$E[\text{Cost to Bob}|\text{Bob lying}] > E[\text{Benefit to Bob}|\text{Bob lying}]$$

We will require

$$\begin{aligned} & 0 > \gamma - d_b \\ & \Rightarrow d_b > \gamma \end{aligned} \quad (5)$$

From equations (4) and (5), we have

$$\gamma < d_b < \frac{p_a c - \gamma_b}{1-p_a} \quad (6)$$

Equations (3) and (6) outline the feasible region for the values of security deposits which will ensure that in sense of expectation rational Alice and Bob have no incentive left to lie. This in turn will make sure that in expectation sense the actual transaction result will be the same as the mediator's declared transaction result. However, it is possible that the feasible region in these equations may not exist because the lower bound for the deposit turns out to be greater than the upper bound. In this case the transaction may not be allowed.

In practice, correct estimation of peer utilities, service costs, and so on, will be a challenging task and will be a limiting factor in success of the security deposit scheme. Another drawback of the security deposit scheme is of a high price paid (as forfeited security deposits) to decrease the overall transaction failure probability. Thus one can argue that a prospective loss of security deposit may increase the "perceived riskiness" of transacting in such a system. To address these issues we propose an "insurance" based payments system in the next section.

	Bob Truthful	Bob Lies
Alice Truthful	$(c + f_a, \gamma_a)$	$(f_a, 0)$
Alice Lies	$(f_a, \gamma_a + \gamma)$	Not possible

TABLE V

(COST,BENEFIT) REPRESENTATION FOR ALICE WITH USE OF INSURANCE

	Bob Truthful	Bob Lies
Alice Truthful	$(\gamma_b, c - f_b)$	$(0, c - f_b + \gamma)$
Alice Lies	$(\gamma_b, c - f_b)$	Not possible

TABLE VI

(COST,BENEFIT) REPRESENTATION FOR BOB WITH USE OF INSURANCE

IV. INSURANCE

This is another approach which can be used to increase robustness against attacks. Herein every transaction between peers (consumer and content owner or content owner and distributor) is insured to an amount of the negotiated price of the transaction. Specifically, the consumer would be interested in seeing that it is refunded the negotiated price of the transaction (which it has pre-payed) in case the transaction fails. Similarly, the provider would like to draw insurance coverage on the payout it was expecting out of a successful transaction which otherwise might be termed as a failure.

A natural choice on which to base the calculation of insurance premiums [10] of these peers would be the estimated probability of truthful behavior and their past performance. One can imagine these premiums being paid over on top of the negotiated price of the transaction. Note that since the mediator chooses to act as the insurer in this scenario, peers can not commit fraud of claiming insurance pay-outs even when the transaction is successful.

It is to be noted that with the use of insurance we can relax the assumption that Alice is always retaliatory as in the case of deposits. This is because with the use of insurance, Alice has incentives to truthfully report the situation in case Bob lies so that she can claim her insurance pay-out.

One can model calculation of insurance premiums in a variety of ways. We provide a sketch of how a simple but functional insurance coverage can be designed for the system in consideration.

The scheme requires that our two transacting peers Alice and Bob pay insurance premiums of say f_a and f_b respectively. If the transaction fails then the insurer has an obligation to refund transaction losses for both peers. The premiums in turn depend on the estimated probabilities of truthful behavior for Alice and Bob, i.e., p_a and p_b. We know that,

$$P(\text{Transaction fails}) = 1 - p_a p_b$$

The content owner will be interested to ensure that in expectation the premiums charged f_a and f_b should be enough to cover the pay-outs due to transaction failure. Thus,

$$f_a + f_b \geq 2c(1 - p_a p_b) \tag{7}$$

But the premiums should not be so large that Alice and Bob do not find any incentives to transact in the system. To ensure that we perform the cost-benefit analysis for both peers.

In the case of Alice, to ensure E[Alice's cost|Alice truthful] < E[Benefit to Alice|Alice truthful] the following should hold

$$(c + f_a)(1 - p_b) + (f_a)p_b < p_b\gamma_a$$

or in other words

$$f_a < p_b\gamma_a - c \tag{8}$$

Similarly, for the case of Bob, the following should hold so that E[Benefit to Bob|Bob truthful] > E[Bob's cost|Bob truthful], we should have

$$(c - f_b)(1 - p_a) + (c - f_b)p_a \geq \gamma_b$$

or,

$$f_b \leq c - \gamma_b \tag{9}$$

It is to be noted that equations (7),(8) and (9) only identify the feasible region for the value of insurance premiums. To determine an exact value of the premium we would need to formulate relationships between premiums and probabilities of truthful behavior and also incorporate formal model of the estimators of the probabilities of truthful behavior with peer's performance history. We consider this as further work.

Also, note that unlike deposits using insurance we can not guarantee that expected cost for a lying peer would be more than her benefit in a single transaction. In-fact as peers would lie often their probabilities of honest behavior would decrease. This in turn would increase the insurance premium resulting in a cost to benefit ratio larger than 1. This will also ensure that the actual and mediator's declared transaction results are same.

Using Insurance to mitigate malicious behavior: If carefully designed, insurance premium of a peer should reflect the potential risk of having that peer conduct transactions in the system. As peers prove themselves to be trust-worthy, their premiums could decrease to nominal levels while the premiums could be substantially higher for peers which demonstrate malicious behavior. Thus it can be used as an effective incentive/disincentive mechanism to regulate peer behavior in the system.

We also believe that insurance is a more sophisticated mechanism than that of security deposits. This is because looking from an individual consumer/distributor point of view this method reduces the 'perceived riskiness' of transacting in the system by providing insurance against transaction failures.

V. CONCLUSION

We presented the problem of designing robust payment schemes for mediated P2P marketplaces with non-monitored transactions. We showed pre-pay-post-pay as the best choice for payment in this setting and identified its weaknesses with regards to a sybil attack. We then proposed supplementing pre-pay-post-pay with deposits or insurance.

We emphasize that this work is of ongoing nature. We do not claim that the techniques outlined in the paper are exhaustive in any sense. On the contrary, we hope that our paper will stimulate further work in the area. In particular we would like to further explore the use of a eBay like reputation system to bolster the pre-pay-post-pay strategy. Another challenging problem is to model the deposit or insurance values using reputations rather than the estimated values of truthful behavior.

Acknowledgements: The authors would like to thank David A. Turner of California State University San Bernadino, Amitabha Bagchi, Ronak Patel and Pulkit Singhal of Polytechnic University for helpful discussions and insightful comments.

References

[1] M. Adler, R. Kumar, K.W. Ross, D. Rubenstein, T. Suel, D. Yao, *"Optimal Peer Selection for P2P Downloading and Streaming,"* IEEE Infocom 2005.

[2] D.A. Turner and K.W. Ross, *"A LightWeight Currency Paradigm for the P2P Resource Market,"* Seventh International Conference on Electronic Commerce Research, Dallas, TX, USA, Jun 2004.

[3] http://www.ebay.com

[4] J.R. Douceur, *"The Sybil Attack,"* Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS02), Cambridge, Massachusetts, USA, March 2002.

[5] http://www.escrow.com

[6] http://www.paypal.com

[7] S. Buchegger and J. Boudec, *"A Robust Reputation System for P2P and Mobile Ad-hoc Networks,"* Second Workshop on Economics of Peer-to-Peer Systems, Cambridge, Massachusetts, USA, June 2004.

[8] Z. Despotovic and K. Aberer, *"Maximum Likelihood Estimation of Peers' Performance in P2P Networks,"* Second Workshop on Economics of Peer-to-Peer Systems, Cambridge, Massachusetts, USA, June 2004.

[9] S.D. Kamvar, M.T. Schlosser, and H. Garcia-Molina, *" The Eigentrust Algorithm for Reputation Mediator in P2P Networks,"* Proceedings of the Twelfth International World Wide Web Conference, Budapest, Hungary, May, 2003.

[10] A.J.G. Cairns, *"A Discussion of Parameter and Model Uncertainty in Insurance,"* Insurance: Mathematics and Economics, 27: 313-330, 2000.