Ptrim: A Market-Based approach to Managing the Risk of Peer-to-Peer Transactions

Stephanos Androutsellis-Theotokis and Diomidis Spinellis

Department of Management Science and Technology,
Athens University of Economics and Business
76 Patission Street, GR-10434, Athens, Greece
{stheotok,dds}@aueb.gr

Abstract. Ptrim is a purely decentralised application for performing peer-to-peer transactions. Instead of relying on a typical trust/reputation management approach, Ptrim allows its users to manage and reduce the risk inherent in their transactions by including a transaction risk management layer on top of the main transaction processing system. This layer is based on the financial principles behind credit derivative markets and the instruments used for transferring and hedging risk therein. Peers can offer to undertake the risk of transactions between other peers, essentially insuring the transactions for the event that they default. A transaction default risk market is thus generated, in which the cost of insuring a transaction reflects the amount of risk involved. Ptrim has been implemented using the JXTA technology, and deployed in a contained environment. We present the Ptrim system design and first trial results. These validate our design decisions and indicate that our system converges towards transaction default insurance costs correlated with the actual transaction risk.

1 Introduction

It is being progressively recognised that information systems and applications supporting collaborative tasks and/or transactions that are based on centralized client-server models, can also be based on the maturing wave of peer-to-peer architectures (see [3]). In order to manage and reduce the risk inherent in peer-to-peer transactions and their decentralised and uncontrolled environment, reputation and trust management systems have been proposed (see Section 5) with the aim to provide peers with estimates of the risk involved in their transactions, based on the observed past behaviour of the collaborating peers.

In this work we present a different approach, based on the financial principles governing credit derivatives markets (see Section 2) for managing, transferring or reducing credit and transaction risk. We target transactions of a financial nature.

Ptrim creates a peer-to-peer derivatives-market-like layer on top of the main transaction processing system, that is used to manage the risk of transaction default. The transacting peers are therefore alleviated from the need to collect,

In *DBISP2P 2007: Proceedings of the Fifth International Workshop on Databases, Information Systems and Peer-to-Peer Computing.* VLDB 2007, Vienna, Austria. September 2007.

process and evaluate reputation and trust information. Instead, they have the option to request offers from other peers to underwrite the risk of their transaction. A transaction default market is thus built on top of the main transaction processing system. The cost of "insuring" a transaction is determined through market-based mechanisms.

We have designed and implemented a system that includes both the main transaction processing tasks, and the creation and operation of the insuring / risk management layer, based on the JXTA platform and technology [21].

In this paper we introduce and discuss the concept of market-based transaction risk management, we describe how we applied this concept to the design and implementation of the Ptrim system, and we present our preliminary deployment scheme, trial results, and the current and future extensions we are working on.

2 Market-based transaction risk management

Our proposed approach to managing transaction risk draws from the general concepts of credit-derivatives markets (see [7, 8, 25]). In finance, derivatives are financial instruments whose performance depends on another variable, and are often used to transfer or reduce (hedge) risk involved in assets or transactions.

We found that a notion that lends itself well to our decentralized architecture and direct one-to-one transaction scenario, is that of a specific category of financial instruments called *Over-The-Counter Credit Default Swaps (CDS)*. CDS's are used to remove risk by directly transferring cash between transacting entities based on privately negotiated contracts, without going through any intermediary (hence their appeal for our peer-to-peer environment). They are most often used as insurance policies, or hedges, to protect against financial loss, and are the most commonly used credit derivatives, corresponding to a total outstanding notional amount of about 300 trillion USD in 2005 [6]. CDS's are usually physically settled upon the occurrence of a *credit default event*. The characterisation of the risk (typically visible as a spread of interest rates) dictates the cost of the derivative products.

2.1 The Ptrim market-based insurance layer

In our design, on top of the main transaction processing system we create another peer-to-peer layer that acts as a market offering transaction default insurance. Any peer can participate in this market: A peer could be regularly active as a buyer or seller, and sporadically choose to also offer insurance for specific transactions (e.g. involving peers that they know can be trusted such as friends, relatives or peers they have interacted with in the past); or another peer could act as a dedicated "insuring institution", only involved in the business of providing insurance for transactions between other peers (for example a bank wishing to obtain additional income from its database of customer credit histories and its IT infrastructure for determining a customer's credit default risk).

The risk default market operates as a derivatives market. The outcomes of insurance transactions are made available to the market by the application, and are utilised to produce new insurance offers. Transacting peers can request offers for insurance protection against their transaction default risks, evaluate such offers and select one (or more, see Section 3) of the most competitive ones. As a result, the transacting peers can significantly reduce their transaction risk, for a small increase in their transaction cost. Furthermore, the cost of the insurance offers received by a peer directly reflects the subjective knowledge of the risk involved (knowledge about past outcomes or lack of such data). This can be used to make a financial decision as to whether to proceed with a transaction or not. On the other hand, peers in the insurance layer who have a picture of the trustworthiness of others, can monetise this knowledge by offering transaction insurance.

The triggering event for settling insurance transactions in our case would be the (unsuccessful) completion of a transaction.

A feature of our proposed approach is therefore that it does not require the peers to be involved in the collection and management of reputation information, or to engage in the decision-making process of whether to proceed with a transaction based on the available reputation information. Note that although very promising work has been done in this direction (see Section 5), in practical terms it can happen that such information may be incomplete, ambiguous, or presented in ways that do not clearly indicate what course of action the peer should follow (e.g. should a transaction involving a $20,000 product be initiated with a peer whose reputation score is 0.67/1.00? What if the product costs $20 instead?).

2.2 Characteristics of our insurance market

In order for the insurance market described above to be efficient, it should approximate to the extent possible the main characteristics of the (largely theoretical) *perfectly competitive markets* [9]:

Atomicity A perfectly competitive market shall comprise a large number of "small" transacting entities whose individual actions can have no impact on others or on the entire market. This is the case for the individual peers in our insurance market and their actions, whose scope is limited per transaction.

Homogeneity There shall be no differentiation in the services offered by the insuring peers. In fact in our market they all offer exactly the same service, so any peer could be substituted by another one.

Perfect and complete information such information is required for the market participants to be able to act rationally. This is ensured by the Ptrim application. Information about the actions of other players in the market is updated upon completion of each transaction event, and circulated by the application across the network of insuring nodes either directly or indirectly following propagated queries (see Section 3).

Equal access & free entry These are guaranteed by Ptrim and the completely decentralized nature of the peer-to-peer architecture it is based upon. All peers function in the same way and are exposed to the same information, and no obstacles are presented to any entity that would wish to enter the network and provide its services.

Individual buyers and sellers act independently This is the characteristic that is notoriously harder to achieve, in any (online or offline) market. Indeed here as well, it is conceivable for a subset of the peers to collude in order to maximize their gain and affect the market price. The extent to which this is possible may not be enough to disturb the main operation of our market, particularly given the availability of complete information about all actions and lack of intermediaries (see [16]), however we have also considered possible countermeasures, as discussed in Section 5.

The above observations show that Ptrim's insurance market layer is an acceptable approximation to a perfectly competitive market. It can thus efficiently support peer-to-peer transactions through the insurance cost offers it provides to the transacting peers.

3 System description

Figure 1 shows a schematic representation of our system and the layers it comprises (transaction processing, risk management, and the underlying infrastructure).

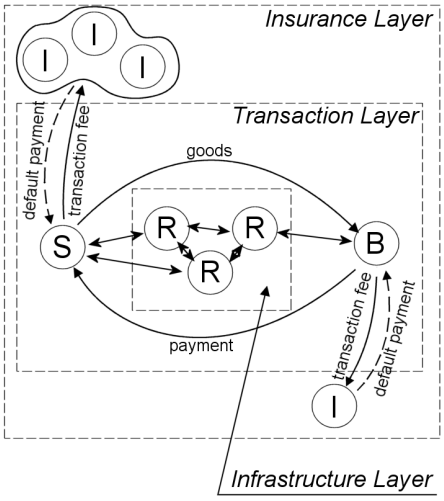

Fig. 1. A schematic abstract representation of the Ptrim system. At the core lies the infrastrucre layer, based on JXTA protocols and a set of dynamic relay peers (R). The main transaction layer consists of buyers (B) and sellers (S) engaged in transactions based on asynchronous message exchanges; above it, the insurance layer includes other peers (I) that offer to insure the transacting peers, either jointly or independently.

The entire system is implemented using the JXTA technology — a set of open protocols that allows connected devices on the network to communicate and collaborate in a peer-to-peer manner. JXTA provides the core communication infrastructure, allowing nodes to enter or leave the network at any time, and messages to be routed around network or node failures, based on a transient number of "rendezvous" (or relay) peers.

Peers can participate in the following three roles:

Rendez-vous peers (R) offer the necessary infrastructure for relaying messages between other peers and providing the network connectivity and availability.

Transacting peers (B,S) advertise products or services they wish to sell, place offers for purchasing other peers' products, and eventually participate in transactions as buyers or sellers.

Insuring peers (I) offer insuring services to the transacting peers, and are remunerated in the form of a percentage of the amount they are insuring.

A peer can participate in one or more of the above functions, and can simultaneously participate in more than one transactions both in the transacting and the insuring roles. For each transaction they participate in, the peers maintain state information as well as records of all exchanged messages and other data.

The peer nodes communicate by exchanging messages asynchronously in a distributed manner. Each message is characterised by a message type and contains an ID and timestamp, sender and recipient identification, as well as other context-specific data fields that are parsed by the recipient (for instance an `OfferForGoods` message will contain identification of the goods advertisement it concerns, financial details, expiration date etc.) The Ptrim documentation lists the types of messages exchanged between peers within the scope of the main transaction processing system and the transaction insurance layer. Most messages are addressed to specified recipients, whereas some are broadcast to entire groups of peers (e.g. `GoodsAdvertisement` messages are broadcast to all transacting peers, and `TransactionInsuranceRequest` to all insuring peers).

The application is available as an open-source project at:
`http://istlab.dmst.aueb.gr/~path/software/ptrim`.

3.1 Main transaction processing system

Within the scope of the main transaction processing system, the transacting peers (or their users) can carry out the following tasks:

- Maintain and update descriptions of products and services that are for sale, and advertise them in the network.
- Place offers for the advertised products that other peers are selling.
- Review offers placed for their advertised products, and choose the most convenient ones. Upon acceptance of an offer, a transaction is initiated between the buyer and the seller. The applications of the two peers initiate local transaction records and exchange messages to bind their respective records with identification data and details, which include deadlines within which they have to be completed.

– Notify the system of the receipt of either the products or the amount due for them, depending on whether the peer is buying or selling, and as a result complete a transaction. Upon receipt of products or payment the applications of the transacting peers exchange messages to bilaterally update their transaction state.
– In case of failed transactions (e.g. if a deadline expired before goods or payment were received), the default event is described and the applications exchange messages notifying the transacting counterpart.

3.2 Transaction insurance layer

Within the scope of the transaction insurance layer, the peers can carry out the following tasks:

– Transacting peers request insurance offers for transactions they are considering engaging in, including descriptions of the product, identification of the transacting counterpart and other details. More than one insuring peers can jointly insure a single transaction. The insurance cost and the potential default payment are fairly distributed among the insuring peers, proportionally to the part of the transaction they are insuring.
– Insuring peers transmit offers for insuring specific transactions.
– Transacting peers review and accept insurance offers, thus initiating insurance transactions. The corresponding insuring peers are notified by the application and respectively initiate insurance transaction records.
– Transacting peers will eventually notify the insuring peers of the transaction outcome, and in the event the transaction defaulted, describe the event and request reimbursement.
– The insuring peers' application will announce to the network information about the outcome of completed insurance transactions. This information is propagate and cached and can be queried by other peer nodes.

3.3 System Operation Example

Figure 2 shows a UML sequence diagram illustrating a simple example of a transaction involving four peers. Peers *p1* and *p2* engage in a transaction as seller and buyer respectively, and peers *p3* and *p4* respond to an insurance request placed by *p1*. Note that this is a relatively straightforward example, with the transaction completing successfully.

Before accepting the offer and initiating the transaction, a transacting peer (in this case the seller (P1)) can broadcast requests for insurance and then waits for offers to be asynchronously received. When enough offers are received to satisfy the request, the transacting peer selects and accepts the best ones (in this case both), and an insurance transaction is initiated. Insurance offers include details such as cost (percentage of transaction cost) and amount to be insured (an insurer has the option to offer to insure only a part of a transaction,

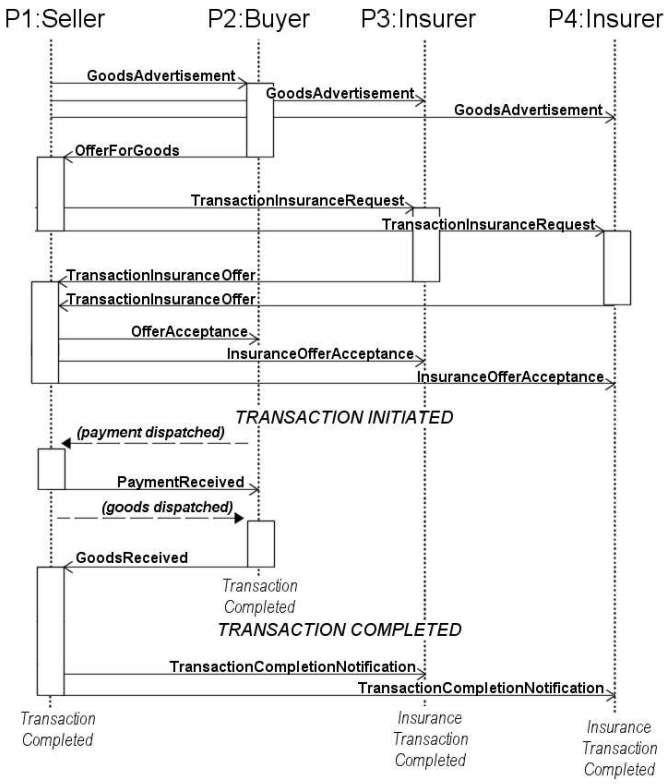

Fig. 2. A UML sequence diagram of a simple transaction

and more than one insurers can jointly insure a single transaction, thus reducing their individual exposure).

When the buyer's offer is accepted, a new transaction record is initiated in both of the transacting peers, and their state is kept synchronized by exchanging messages about the transaction status (whether a product or payment has been dispatched or received, etc).

The transaction is completed upon successful receipt of goods and payment or in case of transaction default, and the insuring peers are notified. Any payments are considered to take place outside the scope of the application.

4 Deployment and preliminary results

Our system has been deployed and is currently in operation within the domain of our academic institution. More wide-scale deployment is scheduled after we enhance it with more robust non-functional characteristics, mainly related to security (see also Section 6).

We carried out sets of small-scale trials within controlled groups of people, using a version of the system specifically targeted towards one-off product trans-

actions. Some members of the group were equipped with products and (fake) money, and others participated as transacting peers. In some cases some of the transacting peers were instructed in confidentiality to exhibit uncooperative behaviour, either refusing to pay for goods purchased, or selling goods of lower quality or specifications than advertised, or otherwise "cheat" in the transactions they were involved in. In order to bootstrap the market, the insuring peers were instructed to initially offer to insure transactions for about 10% of the transaction cost.

The operation of the entire network was followed by a specially designed monitoring peer that regularly requested details about all events and transactions carried out. Table 1 summarises key data of a typical such trial.

Duration:	4 hours
Participating peers:	53 (37 transacting, 16 insuring)
Products advertised:	32
Total messages transmitted or broadcast:	1378
Total transactions:	23 (of which 8 successfully completed)
Average insurance offer rate:	24%

Table 1. Summary of a typical system trial

Due to the small scale of these first trials, no statistically useful quantitative measurements could be collected and presented yet. However, within a short running period and after a few transactions were performed it was observed that, as expected, the peers insuring transactions involving the "cheating" peers began requesting progressively higher rates, as they observed the poor performance of previous insurance transactions involving them. The result was that when peers considered engaging in transactions with cheating peers, they were faced with increasingly more costly insurance offers, which proved to be a reasonable disincentive for them, as they would soon opt out of such transactions and turn to other peers associated with lower insurance costs. The cost of insuring a transaction derived by the insurance market thus proved to be a good measure of the risk involved in a particular transaction.

5 Related work

The concept of market-based transaction risk management presented in this paper is proposed as a potential alternative to distributed reputation management for assisting peers in their transactions. The two however are not mutually exclusive, and a considerable amount of groundbreaking work has been carried out in the distributed reputation management field in recent years. The main aim of this work is to provide an expectation about a peer's behaviour in a transaction, by monitoring, maintaining and distributing information about its behaviour in past transactions.

A variety of solutions have been proposed for addressing either or both of the data modelling, or semantic problem (how to generate, interpret and process the

reputation data), and the data management, or system problem (how to store, retrieve, distribute and secure the reputation data in a scalable and efficient manner) [2]. Some notable systems in this area are the EigenTrust system [22], PeerTrust [30], Credence [28], a system proposed by Aberer et al based on the P-grid structured routing algorithm [2, 1], a Bayesian approach proposed by Buchegger et al [5], TrustMe [26], XRep [11], a partially centralized mechanism presented in [17], to name but a few. Due to space constraints, we refer the reader to a comprehensive overview by Huaizhi and Singhal [19] and the references therein.

Most peer-to-peer networks and reputation management systems are faced with the risk of peers colluding to form malicious collectives or cliques. In our case this is acceptable to a degree, as it reflects the operation of real markets. However a number of approaches to render systems like ours more resistant to collusion have been proposed in the literature and are currently being considered.

In reference [22] the notion of pre-trusted peers is introduced, whereas IP clustering is proposed as a solution in reference [11], but both of these approaches may not be practically applicable in our case.

Among the most interesting approaches, Jurka and Faltings [20] suggest and analyse an incentive-based collusion resistant approach for online reputation mechanisms, while Feldman et al [15] propose the Reciprocative Decision Function, and the use of a maxflow-based reputation management system to achieve optimal levels of cooperation. Zhang et al [31] work on Eigenvector-based reputation systems defining a metric they call Amplification Factor designed to make them more robust to collusion, and in a relevant subject Marti et al [23] propose leveraging the trust associated with social links inherent within peer-to-peer networks for improved DHT routing.

6 Security considerations

In our system description we have not discussed the issue of security, which is very important owing to the nature of the application and the distributed network architecture. We plan to enhance our prototype with provisions to protect if from a variety of security threats, based on solutions proposed in the literature. Table 2 briefly summarises the main security attack categories for applications of this type, and potential countermeasures for each that have been proposed in the literature, and/or already implemented within other peer-to-peer systems and applications.

7 Discussion, future work and conclusions

The Ptrim system proposes an approach to managing the risk of distributed peer-to-peer transactions based on forming and utilising an efficient insurance market. The transacting peers request offers from the insuring peers. The cost of the offers they receive reflects the risk involved, and can be used to make a financial decision as to whether to proceed with a transaction or not. The peers

Attack category	Countermeasure
Data Integrity Attacks (Both for data stored and data routed between nodes)	– Public key infrastructures and encryption (see Credence [29], PeerTrust [30]) – Majority voting mechanisms (see PeerTrust [30]) – Trusted third parties (see EigenTrust [22], TrustMe[26]) – Other cryptographic algorithms/protocols, such as self-certifying data [10], signed files (see Past [14]), information dispersal (see Publius [27], Mnemosyne [18]), secret sharing schemes (see [24])
Data Confidentiality Attacks (both for data stored and routed)	– Public key infrastructures and encryption (see XRep [11], EigenTrust [22]) – Secure routing protocols [10]
Identification-Related Attacks (e.g. pseudospoofing, Sybill attack, ID-stealth, decoy and whitewashing [29], impersonation attacks [12])	– Trusted identity-certifying agents [13] – Public key infrastructures and certificates [29]
Bootstrapping Phase Attacks (e.g. malicious node insertions in the network)	– Majority quorum approaches [12] – Resource-based reputation approaches [11]

Table 2. Main categories of security attacks for peer-to-peer applications, possible countermeasures, and example systems that propose or adopt them. Note that some of the proposed solutions are not purely decentralized (e.g. public key infrastructures)

thus have the option to proceed with a practically risk-free transaction, though with slightly increased cost. At the same time the insuring peers can capitalise on their market knowledge by participating in the transaction insurance market.

We have described our system deployment and first "proof-of-concept" trials, and are now in the process of setting up larger experiments that are expected to produce statistically verifiable results, based on real data collected from the eBay platform. As mentioned in Section 6, the main obstacle to widespread deployment of our system and current work item is the incorporation of necessary security measures.

Another current development plan includes building an API layer around our system in order to incorporate it within our simulation environment (see [4]), and carry out computer-based simulation studies to obtain large data sets for statistical measurements.

We believe that this work may provide an incentive for other researchers to utilise our concept in different applications and cases.

8 Acknowledgements

The authors wish to thank Vassilios Vlachos, George Zouganellis and Sofoklis Stouraitis for their help in setting up the trial environment. This work is implemented within the framework of the Reinforcement Programme of Human Research Manpower (PENED) and co-financed by National and Community Funds.

References

1. Karl Aberer. P-Grid: A self-organizing access structure for P2P information systems. *Lecture Notes in Computer Science*, 2172:179–194, 2001.
2. Karl Aberer and Zoran Despotovic. Managing trust in a peer-2-peer information system. In Henrique Paques, Ling Liu, and David Grossman, editors, *Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM01)*, pages 310–317. ACM Press, 2001.
3. S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys*, 36(4):335–371, December 2004.
4. Stephanos Androutsellis-Theotokis and Diomidis Spinellis. The mor-trust distributed trust management system: Design and simulation results. In *Proceedings of the 2nd International workshop on security and trust management (STM'06)*, Hamburg, Germany, September 2006.
5. Sonja Buchegger and Jean-Yves Le Boudec. A robust reputation system for p2p and mobile ad-hoc networks. In *In Proceedings of the Second Workshop on the Economics of Peer-to-Peer Systems*, 2004.
6. Bank of International Settlements. The BIS Derivative Markets Survey. http://www.bis.org/publ/otc_hy0505.htm, May 2005.
7. AN Bonfim. *Understanding Credit Derivatives and Related Instruments*. Academic Press, December 2004.
8. C Bluhm, L Overbeck, and C Wagner. *An Introduction to Credit Risk Modelling*. CRC Press, September 2002.
9. LMB Cabral. *Introduction to Industrial Organisation*. MIT Press, 2000.
10. M Castro, P Druschel, A Ganesh, Rowstron A, and DS Wallach. Secure routing for structured peer-to-peer overlay networks. In *Proceedings of the 5th Usenix Symposium on Operating Systems*, Boston, MA, December 2002.
11. E Damiani, S De Capitani di Vimercati, S Paraboschi, P Samarati, and F Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the In 9th ACM Conf. on Computer and Communications Security*, Washington DC, November 2002.
12. A Datta, M Hauswirth, and K Aberer. Beyond 'web of trust': Enabling p2p e-commerce. In *Proceedings of the IEEE International Conference on E-Commerce Technology (CEC'03)*, 2003.
13. JR Douceur. The Sybill attack. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)*, MIT Faculty Club, Cambridge, MA, USA, March 2002.
14. P Druschel and A Rowstron. Past: A large-scale, persistent peer-to-peer storage utility. In *Proceedings of the Eighth Workshop on Hot Topics in Operating Systems*, May 2001.

15. M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *Proceedings of the EC04 Conference*, May 2004.
16. Uwe Focht, Andreas Richter, and Jorg Schiller. Intermediation, compensation and collusion in insurance markets. Discussion papers in business administration, Munich School of Management, March 2007.
17. M Gupta, P Judge, and M Ammar. A reputation system for peer-to-peer networks. In *Proceedings of the NOSSDAV'03 Conference*, Monterey, CA, June 1-3 2003.
18. S Hand and T Roscoe. Mnemosyne: Peer-to-peer steganographic storage. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)*, MIT Faculty Club, Cambridge, MA, USA, March 2002.
19. L Huaizhi and M Singhal. Trust management in distributed systems. *IEEE Computer*, 40(2):45–53, February 2007.
20. Radu Jurka and Boi Faltings. Collusion-resistant, incentives-compatible feedback payments. In *Proceedings of the EC07 conference*, San Diego, California, June 2007.
21. The project JXTA web site. http://www.jxta.org, Accessed on-line 2007.
22. Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The EigenTrust algorithm for reputation management in p2p networks. In *Proceedings of the twelfth international conference on World Wide Web*, pages 640–651, New York, 2003. ACM Press.
23. S. Marti, P. Ganesan, and H. Garcia-Molina. DHT routing using social links. In *In Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS 2004).*, 2004.
24. The MojoNation web site. http://www.mojonation.net, Accessed on-line 2007.
25. Robert S. Neal. Credit derivatives: new financial instruments for controlling credit risk. *Economic Review*, (Q II):15–27, 1996.
26. Aameek Singh and Ling Liu. Trustme: Anonymous management of trust relationships in decentralized p2p systems. In *Proceedings of the IEEE Intl. Conf. on Peer-to-Peer Computing*, September 2003.
27. M Waldman, Rubin AD, and Cranor LF. Publius: A robust, tamper-evident, censorship-resistant web publishing system. In *Proceedings of the 9th USENIX Security Symposium*, August 2000.
28. K Walsh and E Gun Sirer. Fighting peer-to-peer spam and decoys with object reputation. In *Proceedings of the SIGCOMM'05 Conference Workshops*, Philadelphia, PA, August 2005.
29. K Walsh and Emin Gon Sirer. Experience with an object reputation system for peer-to-peer filesharing. In *Proceedings of the Symposium on Networked System Design and Implementation*, San Jose, California, may 2006.
30. L Xiong and L Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(16), July 2004.
31. H Zhang, A Goel, et al. Improving eigenvector-based reputation systems against collusion. Technical report, Stanford University, Workshop on Algorithms and Models for the Web Graph (WAW), October 2004.