A Reputation Mechanism for Layered Communities

ELODIE FOURQUET, KATE LARSON and WILLIAM COWAN
David R. Cheriton School of Computer Science, University of Waterloo

The exchange of digital goods via peer-to-peer (P2P) systems is a challenging problem for e-commerce. Participants rarely know each other, and may be completely anonymous, so the self-interest of the participants works against trust and they miss out on the benefits of cooperation. Reputation mechanisms help to remedy selfish misbehaviour. In this paper a new layered reputation mechanism establishes a trusted P2P environment, with bad content filtered out and novel content continuously introduced, by giving appropriate incentives to participants. A simulation was created and experiments run to validate the design.

Categories and Subject Descriptors: K.4.4 [**Computers and Society**]: Electronic Commerce; C.2.4 [**Computer-Communication Networks**]: Distributed System—*Distributed Applications*

General Terms: Algorithm, Design

Additional Key Words and Phrases: Peer-to-peer, reputation mechanism, incentives, collaboration

1. INTRODUCTION

Electronic commerce on the Internet allows strangers to engage in commercial transactions. However, the on-line environment lacks the conventions that provide integrity in face-to-face transactions. For example, villagers know one another, transactions are public and safe. In cities, shared social customs provide enough trust for satisfactory dealings. But online, because anonymity prevails and there are no long-established customs, new mechanisms are needed to provide trust.

Social science research has shown that feedback systems, or *reputation mechanisms*, increase trust and trustworthiness among strangers engaging in commercial transactions [Keser 2003]. They provide summarized histories of past behaviour, increasing the opportunities of well-behaved participants, and decreasing those of poorly-behaved ones. They thus improve trust by rewarding cooperation. In centralized e-commerce settings, such as e-Bay, reputation systems promote trust in anonymous online transactions. Such systems require a trusted central authority to process evaluations of behaviour, which also controls transfers of money and tangible goods between buyers and sellers. Thus, full anonymity rarely exists because the central authority possesses traceable information such as credit card numbers.

Authors' address: E. Fourquet, K. Larson, W. Cowan, David R. Cheriton School of Computer Science, University of Waterloo, 200 University Ave., Waterloo, ON, CANADA N2L 3G1.
All the authors thank NSERC for financial support. One of the authors (EF) thanks Stephen Mann for his advice, encourgement and patience.

P2P transactions differ from centralized ones. Because the goods exchanged are transferred digitally with no money changing hands, the system can provide full anonymity. Unfortunately, full anonymity attracts participants who, noticing the lack of incentive for good behaviour, either exploit the system for their own gain, or weaken it by spreading bad content. Several types of misbehaving participants exist: selfish participants, or *free-riders*, who take resources without giving back, bad and malicious participants who spread poor quality or destructive content, and careless participants who fail to check content quality, passing on good and bad content alike. One consequence is that much content in P2P systems is bad [Liang et al. 2005]. Sanctioning participants who provide bad content and rewarding ones who provide good content is the obvious solution. Without a centralized authority, participants must administer sanctions and rewards themselves. But knowing whom to reward, and when, is a challenging problem. Participants do not know who they can trust; they need a distributed reputation system.

This paper presents a reputation system for P2P which has two design goals: increasing the incentive for good behaviour and improving the quality of content provided. It sanctions participants who provide bad content or who free-ride and rewards those who provide novel, popular content, thus expanding the wealth that the community makes available to its members.

2. RELATED WORK

Peer-to-peer reputation systems require reliable knowledge about transaction partners: without knowledge it is impossible to predict their behaviour, which is essential for interacting. The need first arose in public, centralized systems like eBay. A central trusted server gathers information about every transaction in the system, collates it by participant, and formulates reputation scores which it publicizes. Research has demonstrated that the scores are reliable enough to increase the quantity and quality of transactions [Resnick and Zeckhauser 2002; Dellarocas 2001].

P2P systems, however, lack central trusted authorities. The many reputation systems designed for them gather data in different ways, but all require rigid pseudonyms in order to collate data for each participant. Depending on how the information is stored, network algorithms gather it, labelled by participant, and synthesize reputation scores from it. For example, EigenTrust uses a distributed algorithm to compute global reputation, by collecting information from all participants, requiring rigid pseudonyms, also called weak identity [Kamvar et al. 2003]. Rigid identifiers for collating data are widely regarded as being the *sine quod non* of reputation. For example, 'full anonymity prevents building user reputation', p. 476 of [Marti and Garcia-Molina 2006], or 'In order to be able to trust another peer, a peer needs to know the identity of the other peer', p. 7. of [Suryanarayana and Taylor 2004]. Most important proposals for reputation mechanisms in P2P systems are built on this assumption, which is inappropriate for many P2P systems.

The mechanism described in this paper takes a different approach. It collates information about transactions as the transaction occurs, storing it on the participant whose reputation it affects. This approach assumes the existence of tamper-proof communication and computation, of which much is known. (As one example among many, see [Maña et al. 2004].) On this assumption the mechanism provides a level

of reputation comparable to that of eBay, and without violating anonymity.

Incentive schemes encourage cooperation in P2P networks, decrease selfish behaviour and improve service [Feldman et al. 2004; Feldman and Chuang 2005], but fail sufficiently to combat inappropriate content. Unfortunately, sanction is necessary to reduce bad content: those who spread it must lose their privileges. Most P2P reputation mechanisms facilitate reliable provider selection, reducing a participant's exposure to those who spread bad content, but they do not sanction the disruptive participants spreading bad content [Marti and Garcia-Molina 2006]. The layered mechanism proposed in this paper delivers intrinsic rewards to well-behaved participants, better quality and quantity of content. It curtails inappropriate behaviour by providing only mediocre content to newcomers who free-ride, and removing the privileges of bad participants who spread falsified content and of malicious participants who defect after behaving well. The mechanism rewards provision of new and popular content, guaranteeing wealth for all good participants.

3. LAYERED REPUTATION MECHANISM

Inspired by the concept of "layered communities" [Papaioannou and Stamoulis 2004], the layered reputation mechanism classifies participants into reputation categories which govern their privileges and rewards. Participants in the top layers have a high reputation, having demonstrated good behaviour, participants in the bottom layers have a low, or not yet determined, reputation. A participant's reputation layer determines what content it can access in the system: participants cannot access content more than one layer above them. Thus, participants at the top can access all content, while participants at the bottom have limited access.

Participants can change layers over time only if they supply content to higher ranked participants, who judge the quality of the content received. If the content is judged of good quality by a higher participant the supplier's reputation score increases, possibly moving the supplier to a higher layer. If the content is judged of bad quality by a higher participant the supplier's reputation score decreases, possibly moving the supplier to a lower layer. A key idea is that reputation scores are adjusted non-uniformly: providers of good content in lower layers increase their scores more than those in higher layers. Conversely, supplying poor content from high layers is sanctioned more severely than doing the same from lower ones.

This approach rewards participants who provide popular and good content, moving them to higher layers, which gives them access to better content and isolates them from the demands of free-riders. Conversely, participants who spread bad content fall to lower layers, limiting the harm they produce. They can redeem themselves only by providing good content.

The rest of this section explains the components of the reputation mechanism. First, the representation of participants' reputation information is described. Then the policies which determine content access, judgment process and reputation updates are presented. Finally, the newcomer policy is discussed.

3.1 Participants' Reputations

Every participant has a private *reputation state*, consisting of its reputation score, the amount of good content supplied and the amount of bad, and a public *reputation profile*, consisting of its reputation layer and its badge of honour. A many-to-one

Reputation layer	1	2	3	4	5	6
Range of reputation score	0..15	16..30	31..49	50..70	71..85	86..100
Increment function	9	5	4	3	2	1
Decrement function	1	2	3	4	5	9

Table I. Example of Mechanism Parameters.

function, which protects anonymity, maps participants' reputation scores to their layers. One such function is illustrated in the first two rows of Table I.

The public badge of honour is obtained by a second many-to-one function mapping from the amount of good and bad content supplied to the badge, which gives partial information about the participant's history, while maintaining anonymity. A natural candidate for the function is a discretization of the difference between amounts of good and bad content, but many functions work equally well. The badge of honour neither determines a participant's layer nor effects rewards and punishments, but it plays an important role, weakly signaling trustworthiness, which discriminates among low level participants.

3.2 Content Access Policy

The content accessible to participants is restricted: they may download only from their own layer, the layer above them, and lower layers. For example, a participant in layer four can access content from participants in layers one through five, while a participant in layer one can only access content from layers one and two. Therefore, to acquire content from high layers, participants in low layers must either wait for it to trickle down or move to higher layers. This is the incentive for participants to contribute popular content, the most efficient way of rising. Free-riders, who neither contribute nor rise, have access to limited content.

3.3 Judging Process

When a participant requests content, it receives the reputation profiles of owners from whom downloading is allowed. The requester selects an owner, who then supplies the content. After evaluating it, the requester submits an *appreciation mark*, which updates the owner's amounts of good and bad content supplied. Whether or not the reputation score of the owner changes depends on the layer of the requester.

3.4 Reputation Update Policy

As described above, all receivers of content may submit an appreciation mark, which increments the amounts of good or bad content in the reputation state of the supplier. Appreciation marks from receivers in higher layers also change the reputation score. This restriction prevents receivers in low layers from falsely rating participants in high layers, hoping their content will become available if they fall to lower layers. An exception applies to the highest layer: reputation scores are changed on feedback from receivers in the same layer. There is little incentive for participants in the top layer to lower their neighbours, but if necessary they can sanction participants polluting the top layer.

Reputation scores increase when the appreciation mark is positive, and decrease otherwise. The magnitude of the change depends on the reputation score of the

Participant type	Special	Good	Bad	Malicious
Starting reputation layer	6	2	2	6
Starting reputation score	90	22	22	90
Initial files	No	Yes	Yes	Yes
Check files	Always	Usually	Never	Never
Falsify files	Never	Never	Frequently	Frequently

Table II. Summary of participant behaviour by participant type.

supplier, with different functions for increments and decrements. The increment function decreases with reputation score, while the decrement function increases. Simple examples, the ones used in the simulations, are shown in the third and fourth rows of Table I. They are step functions, with steps occurring at layer boundaries. The layer of the supplier changes when the reputation score changes, if appropriate. These details are critical to the functioning of the mechanism. Participants in high layers must be trustworthy: bad behaviour is punished severely. Good behaviour from participants in lower layers is strongly rewarded in order to encourage it.

3.5 Newcomer Policy

Participants enter the system in a low layer, but preferably not the lowest one. For example, in the simulations, which used a total of six reputation layers, layer two was the entrance layer for newcomers. The content they can access is limited, but if they contribute good and popular content, they quickly rise and access more. This policy discourages whitewashing because a participant who re-enters the system must earn its reputation again. It discourages free-riding because new participants have a strong incentive to contribute novel content. Newcomer requests are often from low layers, and their reputation scores change little. However, high layer requesters are more likely to select newcomers with good badges. Thus, a newcomer who provides good content to its neighbours, invests in future opportunities. However, because all participants contribute to the badge of honour, it is less trustworthy than the participant's layer.

4. EXPERIMENTS

The characteristics of the layered reputation mechanism were examined experimentally using a simulated P2P system. The experiments show the system to have the following desirable properties: participants providing good content rise to higher layers, and ones providing bad content fall; obvious strategies of malicious participants fail to harm the system; participants in high layers have access to better quality content compared to participants in lower layers; and participants that clean downloaded content rise, while participants that falsify content fall.

4.1 Simulation Methodology

4.1.1 *Participant behaviours.* The reputation mechanism has the six reputation layers and the functions defined in Table I. An experiment introduces a population of participants, each participant being one of four types. The starting reputation state, endowment of content, and behaviour of each type are shown in Table II.

Following Kamvar [Kamvar et al. 2003], there initially exist some highly trusted *special participants*. They are the founders of the system, trusted to get the system going. Starting in the top layer with no content, they request content, fetching it initially from the lower layers. They are assumed to evaluate all received content, submitting fair appreciation marks and never falsifying content.

Good participants use the system as intended, introducing good content. As newcomers they start in the entrance layer with an initial endowment of content. They do not falsify content, and usually check the quality of the content they download.

Bad participants spread bad content. Starting as newcomers with an initial endowment of content, they do not check the content they receive, but often falsify it. Good and bad participants anchor a continuum of participant behaviour. *Ordinary participants* check and falsify some of the time, but not all the time: they are better when they check more and falsify less.

Malicious participants deliberately subvert the system, but indirectly. They first attain the highest layer by supplying good content, after which they harm the system by polluting that layer. Assuming they invested prior to the simulation, they start in the highest layer, and behave like bad participants.

4.1.2 *Requesting content and selecting suppliers.* In the simulation the download process and the content downloaded can vary. Downloaded content can vary in popularity, with more popular content more likely to be requested. To implement popularity all content is ranked, with popularity depending inversely on rank, a Zipf distribution, which Gummadi et al. [Gummadi et al. 2003] observed for P2P file requests over unowned files. Varying popularity makes it possible to examine the importance of introducing popular content.

The download process also varies in the simulation. Receiving a list of reputation profiles of owners, the requester may use the profiles to select a supplier. Some participants use random supplier selection, others use best-reputation selection.

4.1.3 *Implementation.* The recommendations of Schlosser et al. [Schlosser et al. 2003] were used to simulate a file-sharing P2P network. Each experiment has a set of participants of several types and an initial endowment of content distributed among the non-special participants. After initialization, the simulation cycle repeats a predetermined number of times to give a run. Fifty runs comprise an experiment. Each cycle of the simulation performs the following steps.

(1) Each participant makes a request, which is successful when the requested content has at least one permissible supplier.
(2) On a successful request, the requester selects a supplier and begins downloading the content. Downloading takes four simulation cycles to finish, which slows the spread of popular files, making the results easier to interpret. An unsuccessful request terminates downloading for the requester on this cycle.
(3) Each participant tests if any downloads completed during the cycle. If so, the content may be evaluated and an appreciation mark submitted. Good content is added to the receiver's endowment; bad content is discarded. Unevaluated content is also added to the receiver's endowment. Bad participants may deliberately falsify content before adding it. When an appreciation mark

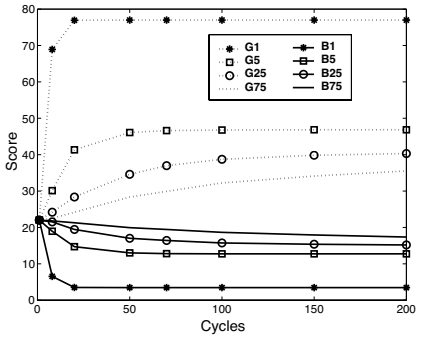

Fig. 1. Variation of the reputation scores of eight participants who provide content varying in quality and popularity. Participants labelled G supply authentic content; ones labelled B supply falsified content. Numbers show the popularity rank of the content.

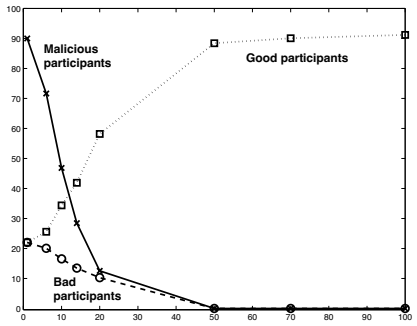

Fig. 2. Variation of reputation scores of three types of participants: good, starting low with good content, bad, starting low with bad content, and malicious, starting high and spreading bad content.

is submitted, the supplier's reputation state and profile are updated.

4.1.4 *Assumptions.* Several assumptions simplify the simulation. First, content is requested and transferred in equal sized atomic units. Second, appreciation marks are not falsified, which may be idealistic, but most marks are provided by well-behaved participants. However, mechanisms that encourage fair rating exist [Papaioannou and Stamoulis 2005] and could be incorporated in the mechanism.

4.2 Results

A pilot experiment examined how quickly newcomers rise as content popularity varies [Fourquet 2005]. It verified the correctness of the Zipf distribution and showed that popularity is important only for the highest popularity content. Less popular content makes only a small difference, which is almost independent of popularity.

4.2.1 *Authentic versus Falsified Content.* The first experiment measures the difference between providing authentic and falsified content, asking how reputation state changes when participants supply good or bad content and how it depends on content popularity.

The system had 10 special participants, plus 4 good and 4 bad newcomers. Good and bad participants each started with one unit of content chosen from a population of 100 units. Good newcomers had authentic and bad newcomers falsified content. Units differed in popularity. Only the special participants requested content, to separate the effect of introducing popular content from the effect of downloading it. Requests for content followed the Zipf distribution with random owner selection, which accentuates the effect of content quality. Each run comprised 200 cycles.

The results are shown in Figure 1. The magnitude of reputation change is monotonic with content popularity, upward for authentic and downward for falsified content. Changes are modest because each newcomer has only a single unit to

share: once authentic content is owned by special participants, they request successfully less frequently. The experiment shows that reputation scores rise when authentic content is supplied, and fall otherwise. Rates of rise and fall depend on the popularity of the content offered.

4.2.2 *Malicious Participants.* The second experiment investigated how much harm malicious participants, who first invest then attack, can cause the system. They aim to maximize harm by attacking from the highest layer. How does this strategy affect the system? Do malicious participants fall faster when attacking than they rose when investing? Is their investment smaller or greater than the harm they cause?

The system had all four types of participants: 10 special, 4 good, 4 bad and 4 malicious. Both special and malicious participants start at the highest layer. Special participants start with no content. Malicious participants falsify all content, trying to spread it to special participants, who select best owners, giving maximal advantage to malicious participants. (Other participants select random owners.) Except for special participants, each participant starts with 20 units chosen randomly from a population of 100 units. Good and special participants always check and never falsify content, whereas bad and malicious participants consistently falsify it. All content is equal in popularity.

Figure 2 shows the results. Malicious participants fall quickly to the entrance layer, falling much more rapidly than participants providing good content rise. The time invested to reach a high reputation score (90) from the entrance point is 4.5 times larger than the time to fall back. As a result, a malicious participant provides almost 5 times more good content than bad content during its rise and subsequent fall. The system actually gets a net benefit from such strategies because of its well-designed increment and decrement functions. Thus, the layered reputation mechanism is resistant to attacks from malicious participants using invest/attack strategies.

4.2.3 *Quality of Content.* The final experiment tests the ability of the layered reputation mechanism to improve content quality. How do differences in falsification and checking rates affect the speed of rising and falling? By how much is the environment better at the top than it is at the bottom?

The simulation includes special, good and bad participants. Ten special participants choose random owners when requesting content. There are 24 good and 24 bad participants, each set divided into 6 groups of 4 participants arranged between good to bad behaviour. Good and bad participants each start with 30 random units of content from a population of 300 units. Good ones check content with various probabilities, and select the best owner at the same probability when requesting content. Bad ones falsify content with various probabilities, always selecting a random owner when requesting. All content has the same popularity, so that all participants receive the same number of requests.

Figure 3 shows the reputation scores of good and bad participants. The more that good participants check content, the higher their reputation scores; the more that bad participants falsify content the lower their scores, as expected. The results have interesting transient behaviour around cycle 75. The reputation scores of

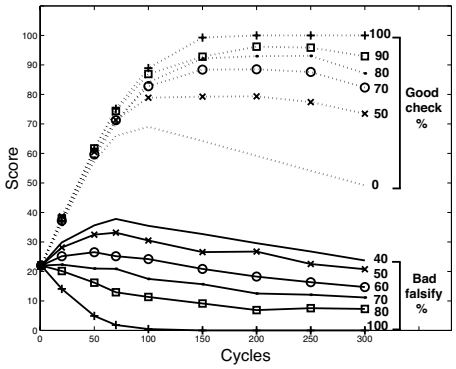

Fig. 3. Variation of reputation score as participants check and falsify content at different rates.

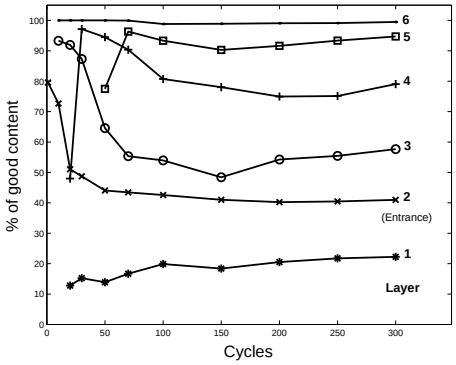

Fig. 4. Proportion of good content in different layers.

good participants who neglect to check (Good 0) initially rise, because they begin with good content. Over time, however, they acquire bad content which they omit to check, and subsequently spread. Their scores then fall. Similarly, bad participants who occasionally falsify (Bad 40, 50) have scores that initially rise because they falsify only some of their good initial content. Soon after, however, they receive bad content which they keep, increasing the proportion of bad content. This effect is aggravated by participants with good content moving up in the layers, their content out of reach of lower participants. Finally, all participants, with the exception of the best (Good 100) eventually decline. Some content exists only in falsified versions, and is increasingly requested as participants come to possess most of the content. High reputation participants, owning one of these falsified versions, experience falling reputations as it is requested. To avoid long term decline, good participants who do not consistently check must, to counteract their negligence, steadily introduce new good content, which benefits the system.

Figure 4 shows the quality of content held by participants in different layers. Three cases explain the data. First, in layer 6 special participants acquire only good content because they check consistently, so the content available in layer 6 is always perfect. Second, from layer 2 good participants move up, taking good, checked content with them, the cause of the peak in content quality that moves from layer 2 into successively higher layers. Third, from layer 2 bad participants move sideways or down. Because they falsify, and do not check, the content quality of the lower layers decreases once the good participants have departed. This robust process shows the substantial reward gained by those exemplary participants who attain layer 6: all possible good content is available in the highest layer, undiluted by falsified content, while the lower layers contain decreasing amounts of good content, increasingly obscured by falsified content.

Figures 3 and 4 thus demonstrate five important properties of the layered mechanism. First, exemplary behaviour (Good 100) is highly rewarded. Second, checking content is essential to maintain a good reputation score, especially when the system contains much falsified content. Third, falsifying content strongly lowers reputation

scores. Fourth, increment and decrement functions, as desired, move participants quickly to the appropriate layer. And fifth, the graded quality of content available in different layers is the reward for behaviour that raises reputation, and the sanction for behaviour that lowers it.

5. FUTURE WORK

This paper describes how decentralized systems of exchange can be improved by a reputation mechanism with built-in incentives. The simulation-based experiments show the overall architecture to function as desired. However, more work can be done so to understand the assumptions made, to investigate mechanism parameters and to test the strength of the mechanism.

First, a concrete distributed implementation of the mechanism is needed. The critical component is a rating daemon in a tamper-proof box on each participant's computer. It accepts appreciation marks, updates the reputation state, and provides the reputation profile. Tamper prevention keeps participants from modifying their own reputation state. To provide appreciation marks, other participants access operations of the daemon. Access can be secured by responsibility trees, which pass encrypted public and private keys to authorized participants. Alternatively, when absolute anonymity is unimportant, central agent processes can update reputation state, which is securely maintained locally [Gupta et al. 2003].

Second, the full effects of system parameters like the number of layers, and the score update functions need to be better understood. Full engineering knowledge of such parameters, which requires large scale simulations, is needed to tune the mechanism for different applications.

Third, malicious strategies must be defined and taxonomized. This paper examined one malicious strategy, invest/attack, which fails to damage the system, but instead improves it. However, other attacks are possible. It seems provable using dynamic optimization [Kamien and Schwartz 1991] that individual attacks fail, but successful collusive attacks are possible. Absolute anonymity makes collusion impossible, and tamper-proof computation of reputation state blocks most techniques for subverting anonymity. However, embedding identity in transferred content remains possible, suggesting that research in what is essentially inverse steganography [Cox et al. 2002] is important. Content transfer architectures like BitTorrent, which disassemble and reassemble content offer suggestive possibilities. BitTorrent is also interesting because it transmits so little bad content [Cohen 2003]. Understanding in human terms why this is so offers another avenue of research.

6. DISCUSSION AND CONCLUSION

A novel mechanism for creating and using reputations was designed specifically for decentralized systems of exchange, meeting the challenge of anonymity and having no recourse to central authority. Without the social context of face-to-face activity, such mechanisms face a variety of problems: free-riding, whitewashing, shortages of novel content, and careless and malicious behaviour.

Productive behaviour is rewarded by better content, so that participants acting selfishly improve the lot of all. Free-riders, for example, are attracted by limited rewards, and become a pool from which active responsible participants are recruited

once they recognize the incentives for introducing new popular content. Careless participants are offered similar incentives to invest in content-checking which improves the system for all. Thus, three problems of participant indolence, free-riding, careless behaviour and content shortage, are solved by the mechanism promising better content in return for genuine contribution. Participants with different trade-offs between the benefits and costs of contribution can co-exist happily in different parts of the system.

Malicious participants and whitewashing are different, challenging the system's defence against deliberate attack. The system defends itself by providing potentially destructive privileges only to participants with a history of contribution. Early detection and quarantine of defecting participants makes the generic malicious attack, invest to gain power then attack from within, counter-productive. The same incentives make whitewashing ineffective, because newcomers have no more power than unsuccessful attackers.

A third issue explored in this work is the tension between anonymity and reputation, which requires participants to know each other. Many-to-one functions mapping reputation state to reputation profile provide enough information for reputation to flourish, but not enough to threaten privacy.

These three ideas are essential components of a successful decentralized community, on-line or real. Recognizing their importance and building them into a community's structure enhances the benefits that all gain from participation. Examples abound of communities that will benefit from layered reputation mechanisms. In coercive environments, users who share controversial thoughts on-line require anonymity to allow free expression and reputation to exclude intrusion. Thus, for example, FreeNet needs reputation. Recent problems at Wikipedia are another example where reputation can ease semi-anonymous electronic collaboration.

Thus, this work offers real promise for evolving decentralized electronic marketplaces, and has the potential to generalize marketplace infrastructure to improve interaction not handled at present by market mechanisms. Only with reputation can electronic communities demonstrate the highly connected efficiency of global economic markets [Seabright 2004].

REFERENCES

Cohen, B. 2003. Incentives build robustness in BitTorrent. In *Proceedings of the 1st Workshop on the Economics of Peer-to-Peer Systems*.

Cox, I., Miller, M. L., and Bloom, J. A. 2002. *Digital Watermarking*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Dellarocas, C. 2001. Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms. In *EC '01: Proceedings of the 3rd ACM Conference on Electronic Commerce*. ACM Press, 171–179.

Feldman, M. and Chuang, J. 2005. Overcoming free-riding behavior in Peer-toPeer systems. *SIGecom Exchanges 5,* 4, 41–50.

Feldman, M., Lai, K., Stoica, I., and Chuang, J. 2004. Robust incentive techniques for Peer-to-Peer networks. In *EC '04: Proceedings of the 5th ACM Conference on Electronic Commerce*. ACM Press, 102–111.

Fourquet, E. 2005. A layered reputation model for Peer-to-Peer system. Class report, CS656, David R. Cheriton School of Computer Science, University of Waterloo.

Gummadi, K. P., Dunn, R. J., Saroiu, S., Gribble, S. D., Levy, H. M., and Zahorjan, J. 2003.

Measurement, modeling, and analysis of a Peer-to-Peer file-sharing workload. In *SOSP '03: Proceedings 19th ACM Symposium on Operating Systems Principles*. ACM Press, 314–329.

Gupta, M., Judge, P., and Ammar, M. 2003. A reputation system for Peer-to-Peer networks. In *NOSSDAV '03: Proceedings 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*. ACM Press, 144–152.

Kamien, M. and Schwartz, N. 1991. *Dynamic optimization: The calculus of variations and optimal control in economics and management*, 2nd ed. Advanced Textbooks in Economics. North Holland, Amsterdam.

Kamvar, S. D., Schlosser, M. T., and Garcia-Molina, H. 2003. The EigenTrust algorithm for reputation management in P2P networks. In *WWW '03: Proceedings 12th International Conference on the World Wide Web*. ACM Press, 640–651.

Keser, C. 2003. Experimental games for the design of reputation management systems. *IBM Systems Journal 42,* 3, 498–506.

Liang, J., Kumar, R., Xi, Y., and Ross, K. 2005. Pollution in P2P file sharing systems. In *INFOCOM '05: Proceedings 24th Annual Joint Conference of the IEEE Computer and Communications Societies*. Vol. 2. 1174–1185.

Maña, A., Lopez, J., Ortega, J. J., Pimentel, E., and Troya, J. M. 2004. A framework for secure execution of software. *International Journal of Information Security 3,* 2, 99–112.

Marti, S. and Garcia-Molina, H. 2006. Taxonomy of trust: Categorizing P2P reputation systems. *Computer Networks 50,* 4, 472–484.

Papaioannou, T. and Stamoulis, G. 2005. An incentives' mechanism for promoting truthful feedback in Peer-to-Peer systems. In *CCGRID '05: Proceedings of 5th IEEE/ACM International Symposium on Cluster Computing and the Grid (Workshop on Global P2P Computing)*. IEEE Computer Society Press, 275–283.

Papaioannou, T. G. and Stamoulis, G. D. 2004. Effective use of reputation in Peer-to-Peer environments. In *CCGRID '04: Proceedings of 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (Workshop on Global P2P Computing)*. IEEE Computer Society Press, 259–268.

Resnick, P. and Zeckhauser, R. 2002. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In *The Economics of the Internet and E-Commerce*, M. R. Baye, Ed. Advances in Applied Microeconomics, vol. 11. Elsevier Science, Amsterdam, 127–157.

Schlosser, M. T., Condie, T. E., and Kamvar, S. D. 2003. Simulating a file-sharing P2P network. In *1st Workshop on Semantics in Grid and P2P Networks*.

Seabright, P. 2004. *The Company of Strangers: A Natural History of Economic Life*. Princeton University Press, Princeton, NJ, USA.

Suryanarayana, G. and Taylor, R. N. 2004. A survey of trust management and resource discovery technologies in Peer-to-Peer applications. Tech. Rep. UCI-ISR-04-6, Institute for Software Research. University of California, Irvine.