

LNAI 4118

Zoran Despotovic  
Sam Joseph  
Claudio Sartori (Eds.)

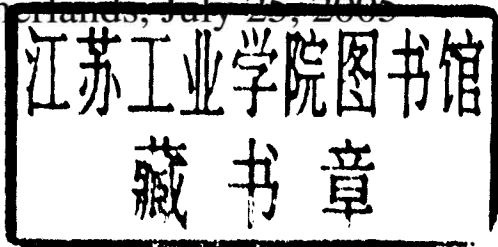
# Agents and Peer-to-Peer Computing

4th International Workshop, AP2PC 2005  
Utrecht, The Netherlands, July 2005  
Revised Papers

Zoran Despotovic Sam Joseph  
Claudio Sartori (Eds.)

# Agents and Peer-to-Peer Computing

4th International Workshop, AP2PC 2005  
Utrecht, The Netherlands, July 25, 2005  
Revised Papers



## Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA  
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

## Volume Editors

Zoran Despotovic  
EPFL Lausanne  
School of Computer and Communication Sciences  
1015 Lausanne, Switzerland  
E-mail: zoran.despotovic@epfl.ch

Sam Joseph  
University of Hawaii  
Dept. of Information and Computer Science  
1680 East-West Road, POST 309, Honolulu, HI 96822, USA  
E-mail: srjoseph@hawaii.edu

Claudio Sartori  
University of Bologna  
Department of Electronics  
Computer Science and Systems  
Viale Risorgimento, 2, 40136 Bologna, Italy  
E-mail: claudio.sartori@unibo.it

Library of Congress Control Number: 2006938336

CR Subject Classification (1998): I.2.11, I.2, C.2.4, C.2, H.4, H.3, K.4.4

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743  
ISBN-10 3-540-49025-6 Springer Berlin Heidelberg New York  
ISBN-13 978-3-540-49025-8 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India  
Printed on acid-free paper SPIN: 11925941 06/3142 5 4 3 2 1 0

# Preface

Peer-to-peer (P2P) computing has attracted enormous media attention, initially spurred by the popularity of file sharing systems such as Napster, Gnutella, and Morpheus. More recently, systems like BitTorrent and eDonkey have continued to sustain that attention. New techniques such as distributed hash-tables (DHTs), semantic routing, and Plaxton Meshes are being combined with traditional concepts such as Hypercubes, Trust Metrics and caching techniques to pool together the untapped computing power at the “edges” of the Internet. These new techniques and possibilities have generated a lot of interest in many industrial organizations, and has resulted in the creation of a P2P working group on standardization in this area (<http://www.irtf.org/charter?gtype=rg&group=p2prg>).

In P2P computing, peers and services forego central coordination and dynamically organize themselves to support knowledge sharing and collaboration, in both cooperative and non-cooperative environments. The success of P2P systems strongly depends on a number of factors. First, the ability to ensure equitable distribution of content and services. Economic and business models which rely on incentive mechanisms to supply contributions to the system are being developed, along with methods for controlling the “free riding” issue. Second, the ability to enforce provision of trusted services. Reputation-based P2P trust management models are becoming a focus of the research community as a viable solution. The trust models must balance both constraints imposed by the environment (e.g., scalability) and the unique properties of trust as a social and psychological phenomenon. Recently, we are also witnessing a move of the P2P paradigm to embrace mobile computing in an attempt to achieve even higher ubiquitousness. The possibility of services related to physical location and the relation with agents in physical proximity could introduce new opportunities and also new technical challenges.

Although researchers working on distributed computing, multi-agent systems, databases and networks have been using similar concepts for a long time, it is only fairly recently that papers motivated by the current P2P paradigm have started appearing in high-quality conferences and workshops. Research in agent systems in particular appears to be most relevant because, since their inception, multi-agent systems have always been thought of as collections of peers.

The multi-agent paradigm can thus be superimposed on the P2P architecture, where agents embody the description of the task environments, the decision-support capabilities, the collective behavior, and the interaction protocols of each peer. The emphasis in this context on decentralization, user autonomy, dynamic growth and other advantages of P2P also leads to significant potential problems. Most prominent among these problems are coordination—the ability of an agent to make decisions on its own actions in the context of activities of other agents—and scalability—the value of the P2P systems lies in how well

they scale along several dimensions, including complexity, heterogeneity of peers, robustness, traffic redistribution, and so forth. It is important to scale up coordination strategies along multiple dimensions to enhance their tractability and viability, and thereby to widen potential application domains. These two problems are common to many large-scale applications. Without coordination, agents may be wasting their efforts, squandering resources and failing to achieve their objectives in situations requiring collective effort.

This workshop brought together researchers working on agent systems and P2P computing with the intention of strengthening this connection. Researchers from other related areas such as distributed systems, networks and database systems were also welcome (and, in our opinion, have a lot to contribute). We seek high-quality and original contributions on the general theme of “Agents and P2P Computing.” The following is a non-exhaustive list of topics of special interest:

- Intelligent agent techniques for P2P computing
- P2P computing techniques for multi-agent systems
- The Semantic Web and semantic coordination mechanisms for P2P systems
- Scalability, coordination, robustness and adaptability in P2P systems
- Self-organization and emergent behavior in P2P systems
- E-commerce and P2P computing
- Participation and contract incentive mechanisms in P2P systems
- Computational models of trust and reputation
- Community of interest building and regulation, and behavioral norms
- Intellectual property rights and legal issues in P2P systems
- P2P architectures
- Scalable data structures for P2P systems
- Services in P2P systems (service definition languages, service discovery, filtering and composition etc.)
- Knowledge discovery and P2P data mining agents
- P2P-oriented information systems
- Information ecosystems and P2P systems
- Security considerations in P2P networks
- Ad-hoc networks and pervasive computing based on P2P architectures and wireless communication devices
- Grid computing solutions based on agents and P2P paradigms
- Legal issues in P2P networks

The workshop series emphasizes discussions about methodologies, models, algorithms and technologies, strengthening the connection between agents and P2P computing. These objectives are accomplished by bringing together researchers and contributions from these two disciplines but also from more traditional areas such as distributed systems, networks, and databases.

This volume is the post-proceedings of AP2PC 2005, the Fourth International Workshop on Agents and P2P Computing,<sup>1</sup> held in Utrecht, Netherlands on

---

<sup>1</sup> <http://p2p.ingce.unibo.it/>

July 25, 2005 in the context of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005).

This volume brings together papers presented at AP2PC 2005, fully revised to incorporate reviewers' comments and discussions at the workshop. The volume is organized according to the following sessions held at the workshop:

- P2P Networks and Search Performance
- Emergent Communities and Social Behaviors
- Semantic Integration
- Mobile P2P systems
- Adaptive Systems
- Agent-Based Resource Discovery
- Trust and Reputation

We would like to thank the invited speaker Hector Anthony Rowstron, Senior Researcher from Microsoft Research in Cambridge UK, for his talk entitled "Removing the Overlay from an Underlay!"

We would also like to thank Omer Rana, from the Department of Computer Science at Cardiff University, UK, for chairing the panel with the theme "To Trust or Not to Trust." We express our deepest appreciation to the workshop participants (more than 40) for their lively discussions, in particular for the invited panelists: Simon Miles, Maria Gini, Martin Purvis and Cristiano Castelfranchi. Many thanks also to Raj Dasgupta and Karen Fullam for chairing sessions in the workshop.

After distributing the call for papers for the workshop, we received 27 papers. All submissions were reviewed for scope and quality, and 13 were accepted as full papers. We would like to thank the authors for their submissions and the members of the Program Committee for reviewing the papers under time pressure and for their support of the workshop. Finally, we would like to acknowledge the Steering Committee for its guidance and encouragement.

This workshop followed the successful third edition, which was held in conjunction with AAMAS in New York City in 2004. In recognition of the interdisciplinary nature of P2P computing, a sister event called the International Workshop on Databases, Information Systems, and P2P Computing<sup>2</sup> was held in Trondheim, Norway in August 2005 in conjunction with the International Conference on Very Large Data Bases (VLDB).

September 2005

Zoran Despotovic  
Sam Joseph  
Claudio Sartori

---

<sup>2</sup> <http://dbisp2p.ingce.unibo.it/>

# Organization

## Executive Committee

### Organizers

Program Co-chairs

Zoran Despotovic  
School of Computer and Communications Sciences,  
Ecole Polytechnique Fédérale de Lausanne (EPFL)  
CH-1015 Lausanne, Switzerland  
E-mail: zoran.despotovic@epfl.ch

Sam Joseph  
Dept. of Information and Computer Science,  
University of Hawaii  
1680 East-West Road, POST 309, Honolulu, HI 96822,  
USA  
E-mail: srjoseph@hawaii.edu

Claudio Sartori  
Dept. of Electronics, Computer Science and Systems,  
University of Bologna  
Viale Risorgimento, 2 - 40136 Bologna, Italy  
E-mail: claudio.sartori@unibo.it

Panel Chair

Omer Rana  
School of Computer Science, Cardiff University  
Queen's Buildings, Newport Road,  
Cardiff CF24 3AA, UK

### Steering Committee

Karl Aberer, EPFL, Lausanne, Switzerland

Sonia Bergamaschi, Dept. of Science Engineering,  
University of Modena and Reggio-Emilia, Italy

Manolis Koubarakis, Dept. of Electronic and Computer Engineering,  
Technical University of Crete, Greece

Paul Marrow, Intelligent Systems Laboratory,  
BTexact Technologies, UK

Gianluca Moro, Dept. of Electronics, Computer Science and Systems, Univ. of Bologna, Cesena, Italy

Aris M. Ouksel, Dept. of Information and Decision Sciences, University of Illinois at Chicago, USA

Claudio Sartori,  
IEIIT-BO-CNR, University of Bologna, Italy

Munindar P. Singh,  
Dept. of Computer Science, North Carolina State University, USA

### **Program Committee**

Karl Aberer, EPFL, Lausanne, Switzerland  
Alessandro Agostini, ITC-IRST, Trento, Italy  
Djamal Benslimane, Universite Claude Bernard, France  
Sonia Bergamaschi, University of Modena and Reggio-Emilia, Italy  
M. Brian Blake, Georgetown University, USA  
Rajkumar Buyya, University of Melbourne, Australia  
Paolo Ciancarini, University of Bologna, Italy  
Costas Courcoubetis, Athens University of Economics and Business, Greece  
Yogesh Deshpande, University of Western Sydney, Australia  
Asuman Dogac, Middle East Technical University, Turkey  
Boi V. Faltings, EPFL, Lausanne, Switzerland  
Maria Gini, University of Minnesota, USA  
Dina Q. Goldin, University of Connecticut, USA  
Chihab Hanachi, University of Toulouse, France  
Willem Jonker, Philips, Netherlands  
Mark Klein, Massachusetts Institute of Technology, USA  
Matthias Klusch, DFKI, Saarbrücken, Germany  
Tan Kian Lee, National University of Singapore, Singapore  
Zakaria Maamar, Zayed University, UAE  
Wolfgang Mayer, University of South Australia, Australia  
Dejan Milojicic, Hewlett Packard Labs, USA  
Alberto Montresor, University of Bologna, Italy  
Luc Moreau, University of Southampton, UK  
Jean-Henry Morin, University of Geneva, Switzerland  
Andrea Omicini, University of Bologna, Italy  
Maria Orlowska, University of Queensland, Australia  
Aris. M. Ouksel, University of Illinois at Chicago, USA  
Mike Papazoglou, Tilburg University, Netherlands  
Paolo Petta, Austrian Research Institute for AI, Austria,  
Jeremy Pitt, Imperial College, UK  
Dimitris Plexousakis, Institute of Computer Science, FORTH, Greece



Martin Purvis, University of Otago, New Zealand  
 Omer F. Rana, Cardiff University, UK  
 Douglas S. Reeves, North Carolina State University, USA  
 Thomas Risse, Fraunhofer IPSI, Darmstadt, Germany  
 Pierangela Samarati, University of Milan, Italy  
 Heng Tao SHEN, ITEE, UQ, Australia  
 Christophe Silbertin-Blanc, University of Toulouse, France  
 Maarten van Steen, Vrije Universiteit, Netherlands  
 Katia Sycara, Robotics Institute, Carnegie Mellon University, USA  
 Peter Triantafyllou, Technical University of Crete, Greece  
 Anand Tripathi, University of Minnesota, USA  
 Vijay K. Vaishnavi, Georgia State University, USA  
 Francisco Valverde-Albacete, Universidad Carlos III de Madrid, Spain  
 Maurizio Vincini, University of Modena and Reggio-Emilia, Italy  
 Fang Wang, BTexact Technologies, UK  
 Gerhard Weiss, Technische Universität, München, Germany  
 Bin Yu, North Carolina State University, USA  
 Franco Zambonelli, University of Modena and Reggio-Emilia, Italy

### **Preceding Editions of AP2PC**

Here are the references to the preceding editions of AP2PC, including the volumes of revised and invited papers:

- AP2PC 2002 was held in Bologna, Italy, July 15, 2002. The Web site can be found at <http://p2p.ingce.unibo.it/2002/> The proceedings were published by Springer as LNCS volume no. 2530 and are available online at: <http://www.springerlink.com/content/978-3-540-40538-2/>
- AP2PC 2003 was held in Melbourne, Australia, July 14, 2003. The Web site can be found at <http://p2p.ingce.unibo.it/2003/> The proceedings were published by Springer as LNCS volume no. 2872 and are available online at: <http://www.springerlink.com/content/978-3-540-24053-2/>
- AP2PC 2004 was held in New York City, USA, July 19, 2004. The Web site can be found at <http://p2p.ingce.unibo.it/2004/> The proceedings were published by Springer as LNCS volume no. 3601 and are available online at: <http://www.springerlink.com/content/978-3-540-29755-0/>

# Table of Contents

## Trust and Reputation

Optimizing an Incentives' Mechanism for Truthful Feedback in Virtual Communities .....	1
<i>Thanasis G. Papaioannou, George D. Stamoulis</i>	
A New View on Normativeness in Distributed Reputation Systems: Beyond Behavioral Beliefs .....	16
<i>Philipp Obreiter, Birgitta König-Ries</i>	
A Trust Management Scheme in Structured P2P Systems .....	30
<i>So Young Lee, O-Hoon Kwon, Jong Kim, Sung Je Hong</i>	
Incentive-Compatibility in a Distributed Autonomous Currency System .....	44
<i>Kenji Saito, Eiichi Morino, Jun Murai</i>	
Handling Free Riders in Peer-to-Peer Systems .....	58
<i>Loubna Mekouar, Youssef Iraqi, Raouf Boutaba</i>	

## P2P Infrastructure

Highly Available DHTs: Keeping Data Consistency After Updates .....	70
<i>Predrag Knežević, Andreas Wombacher, Thomas Risse</i>	
Caching Indices for Efficient Lookup in Structured Overlay Networks ....	81
<i>Vasilios Darlagiannis, Nicolas Liebau, Oliver Heckmann, Andreas Mauthe, Ralf Steinmetz</i>	

## Semantic Infrastructure

A Semantic Marketplace of Negotiating Agents .....	94
<i>Theodore Patkos, Dimitris Plexousakis</i>	
Semantic Web Service Composition Through a P2P-Based Multi-agent Environment .....	106
<i>Peep Küngas, Mihhail Matskin</i>	

Community and Mobile Applications

A Low-Latency Peer-to-Peer Approach for Massively Multiplayer Games..... 120  
    *Jin Zhou, Li Tang, Kai Li, Hao Wang, Zhizhi Zhou*

An Agent-Based Collaborative Framework for Mobile P2P Applications..... 132  
    *Menggiu Wang, Heiko Wolf, Martin Purvis, Maryam Purvis*

ACP2P: Agent-Community-Based Peer-to-Peer Information Retrieval – An Evaluation..... 145  
    *Tsunenori Mine, Akihiro Kogo, Makoto Amamiya*

A Peer Ubiquitous Multi-agent Framework for Providing Nomadic Users with Adapted Information ..... 159  
    *Angela Carrillo Ramos, Jérôme Gensel, Marlène Villanova-Oliver, Hervé Martin*

Author Index ..... 173

# Optimizing an Incentives' Mechanism for Truthful Feedback in Virtual Communities<sup>\*</sup>

Thanasis G. Papaioannou and George D. Stamoulis

Department of Informatics, Athens University of Economics and Business (AUEB)  
76 Patision Str., 10434 Athens, Greece  
{pathan, gstamoul}@aueb.gr

**Abstract.** We analyze a mechanism that provides strong incentives for the submission of truthful feedback in virtual communities where services are exchanged on a peer-to-peer basis. Lying peers are punished with a severity that is exponential to their frequency of lying. We had first introduced and evaluated experimentally the mechanism in [1]. In this paper, we develop a Markov-chain model of the mechanism. Based on this, we prove that, when the mechanism is employed, the system evolves to a beneficial steady-state operation even in the case of a dynamically renewed population. Furthermore, we develop a procedure for the efficient selection of the parameters of the mechanism for any peer-to-peer system; this procedure is based on ergodic arguments. Simulation experiments reveal that the procedure is indeed accurate, as well as effective regarding the incentives provided to participants for submitting truthful feedback.

## 1 Introduction

Virtual communities for the exchange of files, services, knowledge or opinions possibly on a peer-to-peer basis have already been widely developed. In the absence of any proper accounting about who is offering value to others in such communities, there is opportunity for free-riding and for malicious actions against other members. Revelation of hidden information on the quality of the exchanged good and on the trustworthiness of the community members is necessary. For, otherwise, such virtual environments may offer low value and eventually collapse. Reputation on the basis of ratings can be a proper means for achieving accountability. However, reputation mechanisms are vulnerable to *false* or *strategic voting (rating)*. For example, a particular peer may benefit by submitting unjustified positive ratings for his friends or his collaborators, and/or by submitting unfair negative ratings for his competitors. This problem is further augmented in case of pseudo-spoofing, i.e. use of multiple false identities, which may arise in virtual environments, especially peer-to-peer systems. In [1], we proposed a mechanism for providing incentives for credible reporting of feedback information in a peer-to-peer system. The mechanism was combined with

---

<sup>\*</sup> The present work was partly funded by the IST project EuroNGI (IST-2003-507613).

reputation-based policies that we introduced in [2]. These determine the pairs of peers that are eligible to transact, in order incentives to peers for offering better services to others to be provided as well. According to the mechanism both transacting peers (rather than just the client) submit ratings on the performance of their mutual transaction. If these ratings are in *disagreement*, then *both* transacting peers are punished, since such an occasion is a sign that one of them is lying, yet the system cannot tell which one. When under punishment, a peer is not allowed to transact with others. The severity (i.e. duration) of each peer's punishment is determined by his corresponding non-credibility metric; this is maintained by the mechanism and evolves according to the peer's record. Simulation experiments in [1] showed clearly that the combination of the mechanism with reputation-based policies detects and isolates liar peers effectively, while rendering lying costly even in dynamically evolving peer-to-peer systems. Also, the efficiency losses induced to sincere peers by the presence of large subsets of the population of peers that provide their ratings either falsely or according to various unfair strategies are diminished. As explained in [1], this mechanism can be implemented in practical cases of peer-to-peer systems.

In this paper, we analytically study the standalone effectiveness of the mechanism of [1] (i.e. without being combined with reputation-based policies) in providing incentives for truthful reporting. We define a Markov-chain model in order to study the steady-state effect of the credibility mechanism in punishing liar peers. We also develop an optimization procedure for the determination of the proper parameters of the credibility mechanism employed to a dynamically renewed peer-to-peer system, so as to maximize the effectiveness of the mechanism in punishing lying and minimize the cost induced to sincere peers by potential unfair punishments thereof due to the mechanism. This optimization procedure is based on ergodic arguments. We evaluate our Markovian model and our optimization procedure by simulation experiments that show the accuracy and the effectiveness of the approach. The results scale for realistic population sizes of peer-to-peer systems thus making both our mechanism and our approach for selecting its parameters applicable in practical cases.

There is significant related work in the literature. Dellarocas deals in [3] with the problem of unfair ratings and discriminatory behavior in on-line trading communities. Schillo *et al.* [4] deal separately with behavior and credibility of other agents using the so-called disclosed prisoners' dilemma game with partner selection based on own observations. Damiani *et al.*, in a similar approach [5], extend Gnutella protocol to calculate performance and credibility of other peers based on a peer's own experience and on votes from witnesses. A single trust metric is used for credibility and performance by Yu *et al.* in [6]. Aberer *et al.* [7] present an approach to evaluate trustworthiness (i.e. the combination of credibility and performance) of peers based on the complaints posed for them by other peers following transactions. An approach for providing incentives for truthful reporting of feedback in e-markets has been proposed by Jurca and Faltings in [8]. This approach, similarly to ours, employs disagreement in feedback messages for discovering potential lying. Detailed comparison of our credibility

mechanism with these works has been done in [1]. However, these approaches (including [1]) mostly resort to simulation for the purpose of evaluation of their mechanisms. Moreover, they do not deal with large fractions of collaborated liar peers, as opposed to both [1] and the present work.

The remainder of this paper is organized as follows: in Section 2, we overview our credibility mechanism. In Section 3, we describe the Markov-chain model of a peer-to-peer system that employs our credibility mechanism. In Section 4, we present our procedure for the optimization of the parameters of the credibility mechanism for a peer-to-peer system. In Section 5, we evaluate our Markov-chain model and our optimization procedure by simulation experiments. Finally, in Section 6, we provide some concluding remarks.

## 2 The Credibility Mechanism

Consider a peer-to-peer system for exchanging services that employs a distributed reputation system for performance. Time is assumed to be slotted. For simplicity, we assume that the minimum time interval between two successive service requests by the same peer equals one time slot. Following a transaction, the client peer sends feedback rating his offered performance. For example, he may rate the transaction as "successful" (i.e. high offered performance) or as "unsuccessful" (i.e. low offered performance). The feedback messages are useful only if their content is *true*. Unfortunately, peers actually have the incentive of strategic rating of others' performance, since they can thus hide their poor performance, improve their reputation, and possibly take advantage of others. Thus, a proper mechanism should make lying costly or at least unprofitable. "Punishing liars" is a known recipe [9], [10], but two questions arise: How can lying peers be discovered? How can they be punished in a peer-to-peer system, where there is no central control?

Under our approach peers submit ratings' feedback according to the following rules: i) after a transaction, *both* peers involved have to send one feedback message each, and ii) besides rating (i.e. voting) the transaction as successful or not, each feedback message *also* contains a quantifiable performance metric, e.g. the number of transferred bytes of useful content. We assume that the observed performance is with high probability the same with that actually offered. (The opposite may only occur due to unexpected events during a transaction like network congestion etc.) Thus, if feedback messages for a transaction *disagree* (either in their performance metric or in their vote), then, with high probability, at least one of the transacted peers is lying and has to be somehow *punished*, in order for the right incentives to be provided. However, the system cannot tell which of the peers does lie, and consequently whom to believe and whom to punish. Thus, according to our approach, *both* peers are punished in this case. This idea was initially introduced in [9]. However, by simply applying it, a sincere peer is often punished unfairly.

Therefore, we need a complete mechanism specifying how to punish peers in such an uncontrolled system and how to limit potential unfairness. To this end,

we introduce for each peer: i) the *non-credibility* metric  $ncr$ , which corresponds to reputation for non-credibility, and ii) a binary *punishment state* variable, declaring whether the peer is "under punishment" (if the variable is "true") or not (if the variable is "false"). For each peer, both  $ncr$  and punishment state are public information, and they are appropriately stored so that they are available to other peers. (See [1] for a discussion on practical implementation.) Upon entering the peer-to-peer system, each peer is assigned a positive non-credibility value  $ncr_0$ , while he is not under punishment. (Note that the lower the value of  $ncr$  the better.) This choice of  $ncr_0$  limits the incentive for name changes after a disagreement. The flowchart of the credibility mechanism is depicted in Figure 1. In particular, after a transaction between two not punished peers  $i, j$  their feedback messages  $f_i, f_j$  are sent as input to the mechanism: Upon *disagreement* (i.e. if  $f_i \neq f_j$ ), the non-credibility values of the transacted peers are both increased by  $x$  while both get punished. The duration of a peer's punishment equals  $b_{ncr}$ , i.e. is exponential in his non-credibility, with a base  $b > 1$ . Upon *agreement* (i.e. if  $f_i = f_j$ ), the non-credibility values of the transacted peers are decreased (i.e. improved) by  $d$ , where  $0 < d \leq x$ , without ever dropping below 0. In the rest of the paper, without loss of generality, we take  $x = 1$ . The common feedback is forwarded to the system computing reputation for performance.

Decrease of non-credibility in cases of agreement serves as a rehabilitation mechanism. This is crucial for the efficient operation of the credibility mechanism, because, as already mentioned, upon disagreement in reports, most probably one peer is unfairly punished. The value of  $d$  determines the speed of restoring a non-credible reporting behavior. We employ additive increase/decrease of the non-credibility values for simplicity. Other approaches such as additive increase/multiplicative decrease are also possible.

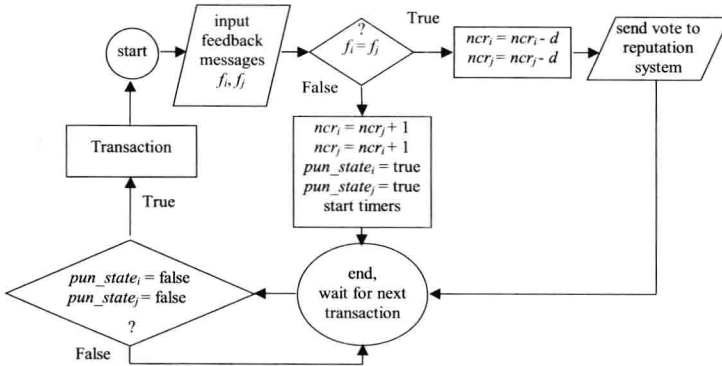


Fig. 1. The credibility mechanism

Punishing peers is not an easy task to employ in the absence of any control mechanism, particularly if peers have full control over their part of peer-to-peer middleware. In our mechanism, a punishment amounts to loss of value offered

by other peers. That is, a peer under punishment does *not* transact with others during his punishment period, while his ratings for such transactions are *not* taken into account. The latter measure provides incentives for peers to abide with the former one! Indeed, first, note that sincere peers under punishment are not expected to be willing to offer services as they would be subject to strategic voting without being able to disagree. On the other hand, liar punished peers collaborated with other liar peers that strategically vote them (i.e. always positively) can raise their reputation anyway, thus having no incentives to perform well during their punishment. Thus, no peer has any incentives to *ask* for services from a punished peer except for strategic voting. Moreover, no peer has any incentive to perform well when offering services to a punished peer, because the corresponding feedback is not taken into account. Therefore, it is beneficial for the system to *prohibit* transaction with punished peers by rule. To this end, if a peer transacts with a punished one, then both of the transacting peers are punished as if they were involved in a new disagreement. Thus, the non-credibility value of a peer remains unchanged during his punishment period unless he transacts with other peers; in such a case it is further increased.

Peers should have the incentive to submit feedback, despite the risk of disagreement and subsequent punishment. Indeed, after a transaction that failed peers may not be willing to report the failure at all. Thus, to provide peers with the incentive to submit their feedback, our mechanism punishes both peers involved in a transaction if only one of them submits feedback. This also prevents unilateral submission of feedback messages for non-existing transactions. Note also that, since the proposed mechanism improves the long-term efficiency of the sincere peers, only liar peers are expected to have incentives to avoid submitting feedback. Yet applying the reasoning of [11] to our case, we expect that under certain circumstances, the existence of our mechanism will lead liar peers to give up their strategic behavior since it is not beneficial to them.

### 3 The Markovian Model Approximating the Mechanism

In this section, we analytically study the *effectiveness* of the proposed mechanism in *equilibrium* for providing incentives to peers for truthful reporting. For this purpose, we define a discrete-time Markov-chain model of a peer-to-peer system where the credibility mechanism is employed. Then, we derive the *steady-state* distribution of the punishment state of sincere and liar peers of the modeled peer-to-peer system. Modeling of time is different than that introduced in Section 2. In particular, for the purpose of specifying and analyzing this Markov chain, we define as time step of our discrete-time model the interval between two successive service requests. We assume that in this interval at most *one* transaction takes place. Thus, transition from one state to another can *only* happen after a transaction between two peers. This is very convenient for analyzing the Markov-chain model and studying the performance of the original system defined in Section 2. Performance measures can be easily translated from the new "transaction units" to actual time slots; see Section 4. Note that at the



beginning of each time step, a peer is randomly selected to be the client of the only transaction that takes place in this step.

The total populations of sincere and liar peers in the peer-to-peer system modelled as a Markov chain are  $S_0$  and  $L_0$  respectively. Consider that a state is a snapshot of the system where state variables are the number of not punished sincere peers  $s$ , the number of not punished liar peers  $l$ , and the number of peers under punishment  $k$ . Clearly, this Markov chain has  $(S_0 + 1)(L_0 + 1)$  different states. Observe also that state variable  $k$  can be computed by the formula  $k = S_0 - s + L_0 - l$ , but  $k$  is used for readability reasons. Let  $q$  be the probability that a requested service is found at a certain peer and  $r$  to be the probability that a peer asks for a service. Recall that credibility values and punishment state are *public* information, and that not punished peers are not allowed to transact with punished peers. The probability that a selected client peer finds a requested service is given by:

$$y = r(1 - (1 - q)^{l+s-1}) \quad (1)$$

A client sincere peer is punished if he finds his service at a liar peer. The probability PS of this event is given by:

$$P_S = \frac{l}{s + l - 1} y \quad (2)$$

A client liar peer is punished if he interacts with a sincere peer plus or with another liar peer that is not collaborated with. We assume that the probability of each given pair of liars to be collaborated with each other equals  $a$ , which is fixed. Thus, the probability of punishment for a client liar peer is given by the formula below:

$$P_L = \frac{s}{s + l - 1} y + \frac{l - 1}{s + l - 1} ya \quad (3)$$

If no liars are collaborated with each other, then  $a = 1$ , while for all liars being collaborated with each other  $a = 0$ . In the analysis that follows we study the case where all liar peers are collaborated with each other, which is the hardest one for the mechanism to deal with.

Recall that at the beginning of each time step, a peer is randomly selected to be the client of the only transaction to take place. The probability  $P_T$  that the two peers of a transaction are punished, i.e. they disagree in their feedback messages is given by:

$$P_T = \frac{s}{s + l} P_S + \frac{l}{s + l} P_L \quad (4)$$

For modeling purposes, we assume that during a time step, a sincere (resp. liar) peer that is under punishment can be "rehabilitated", i.e. stop being under punishment in the next step, with probability  $P_{RHS}$  (resp.  $P_{RHL}$ ). Thus, when there are  $k = S_0 - s + L_0 - l$  peers under punishment in the current state, the average number of rehabilitated peers in the next state is  $(S_0 - s)P_{RHS} + (L_0 - l)P_{RHL}$ . Next, we relate the Markovian model with the original mechanism of Section 2.