



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2009

Incentives for Voting-based Quality Control and Document Storage in P2P Collaboration Systems

Bocek, T ; Hecht, F V ; Hausheer, D ; Peric, D ; Stiller, B

DOI: <https://doi.org/10.1109/WETICE.2009.43>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-19505>

Conference or Workshop Item

Originally published at:

Bocek, T; Hecht, F V; Hausheer, D; Peric, D; Stiller, B (2009). Incentives for Voting-based Quality Control and Document Storage in P2P Collaboration Systems. In: 18th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE 2009), Groningen, The Netherlands, 29 June 2009. IEEE Computer Society Press, 146-147.

DOI: <https://doi.org/10.1109/WETICE.2009.43>

Incentives for Voting-based Quality Control and Document Storage in P2P Collaboration Systems

Thomas Bocek, Fabio Victora Hecht, David Hausheer, Dalibor Peric, Burkhard Stiller
Communication Systems Group, Department of Informatics IFI, University of Zurich, Switzerland
E-mail: [bocek|hecht|hausheer|peric|stiller]@ifi.uzh.ch

Abstract—Peer-to-peer (P2P) systems achieve scalability, fault tolerance, and load balancing, while effectively reducing infrastructure cost compared to client/server systems. For large wiki-based collaboration systems these characteristics can be beneficial as well. However, a fully decentralized P2P-based collaboration system requires appropriate incentives for users to (a) contribute resources for storing and providing documents, and (b) to participate in collaborative control mechanisms for keeping content quality at a high level. Therefore, this position paper proposes a score-based incentive scheme for voting-based quality control and document storage in a P2P-based collaboration system. The approach is based on the PeerVote mechanism, which enables users to vote on content modifications. In order to encourage peers to store and vote on documents, each peer is assigned a score related to its contribution. Peers have an incentive to achieve a high score which allows them to create and approve new documents. The scheme proposed is resistant to fraud, since the score can be verified by any other peers.

I. INTRODUCTION

Large wiki-based collaboration systems may benefit from the characteristics of peer-to-peer (P2P) systems, which inherently support redundancy, scalability, fault tolerance, and load balancing at a lower cost than a client/server-based system. Several approaches for P2P wiki-based collaboration systems have been proposed ([2], [6], and [7]). In these systems users jointly host, distribute, and modify articles.

However, a key challenge in P2P collaboration systems is to encourage peers to contribute to the system in terms of (a) storage resources, and (b) the participation in collaborative control mechanisms to maintain article quality after each modification. Since article quality depends on a subjective judgment, automatic processing is not possible. P2P reputation and recommendation systems, which are successfully used for entire articles [1], are not feasible for time-critical modifications and different versions of an article, because all modifications require reviews from experts before publishing to prevent arbitrary changes. The fully decentralized, user-based voting scheme PeerVote [3] provides a method to maintain quality of content, as it requires the approval from the majority of previous editors for each modification to a document. However, the provision of appropriate incentives for users to participate and to behave correctly in such a voting scheme remains a challenge.

Thus, a fully-decentralized, score-based incentive scheme is proposed in this position paper. A peer's score is calculated based on its overall contribution to the system, which is measured in terms of storage resources as well as its participation in the PeerVote-based content quality control scheme. The scheme is resistant to fraud, since the score can be recalculated by verifying the contribution of a peer. Verification can be exhaustive or sample-based to limit network overhead, if many verifications are required. To the best of our knowledge, this is the first score-based incentive scheme for a decentralized collaboration system.

The remainder of this position paper is organized as follows: Section II summarizes the PeerVote concept, while Section III presents the design of the score-based incentive scheme. Finally, Section IV concludes the paper.

II. PEERVOTE BACKGROUND

The key concept of PeerVote is that each modification of documents contains meta data of voting results of previous authors, including signed results of each vote. Thus, for proposing a modification, previous editors are requested to review the modification. Only if a majority of these previous editors signs the request within a specified time, the modification is accepted. An example use case of PeerVote is a user, who finds a wrong statement in an article and corrects it. This starts the voting process and requests previous editors to review the change. If the majority of previous authors agrees on the correction, the modified document is published.

PeerVote has been implemented, simulated, and experiments have been run on top of EmanicsLab (cf. [3], [4]). To evaluate PeerVote, a number of articles have been stored in the network and peers proposed modifications to these articles. A modification proposal was either correct or malicious. Malicious peers either voted (a) randomly or (b) in the opposite way, *i.e.* in favor of malicious proposals and against correct proposals. Preliminary evaluation results show that PeerVote is scalable and robust with the presence of malicious peers. Fig. 1 shows the difference between malicious (opposite) voting and random voting as the share of malicious peers is increased from 0% to 50%, while the total number of peers is kept constant at 500. As expected, the share of correct documents drops faster with malicious votes, however, it stays at a high level even with up to 50% malicious peers, partially due to new documents published being always correct.

However, new documents and proposals with few previous authors are prone to malicious effects. With few or no previous authors, proposals can be blocked, if peers do not respond to a review request or always reject the request. A first counter-measure is to use a time-based threshold to lower the threshold limit for stalled documents. A second counter-measure is to introduce an incentive scheme to encourage peers to act correctly, as discussed in the following.




Fig. 1. Effect on documents with malicious and randomly voting peers [3].

III. SCORE-BASED INCENTIVE SCHEME DESIGN

A score-based system provides a moral incentive [5] to a user to increase its score. With such a moral incentive, a user can expect approval or admiration for a high score from its community for every desirable action. The incentive scheme proposed is to increase a peer's score s for each desirable action, which include a successful modification m , vote v , and storage of document d . Accordingly, the score of peer i is defined as follows: $s_i = w_m \cdot n_{m,i} + w_v \cdot n_{v[x],i} + w_d \cdot n_{d,i}$, where $n_{m,i}$ and $n_{d,i}$ denote the number of successful modifications and stored documents of peer i , $n_{v[x],i}$ denotes the number of votes of peer i during the last x voting sessions, and w_m , w_v , and w_d denote the weight of the corresponding actions, which is set by the application provider.

A modification is successful, if previous editors approve it. Approved modifications are stored as meta data. Storing unapproved modifications would lead to a security problem, since malicious peers may cause the meta data to grow large. To avoid this problem, only peers with a high score $s_i \geq s_{High}$ are able to propose a modification. Presenting a modification proposal requires additional information, such as explanations or examples to make the proposal easier to understand. Successful votes are not considered in the score as this would provide an incentive to always follow the majority.

The number of modifications n_m and votes $n_{v[x]}$ can be calculated using the meta data of a document. Since the meta data contains the document creator and all previous voters, a peer only needs to announce which documents to look at for calculating its score. *Previous voters* include only the newest voters, as outlined in [3]. To verify the score of a peer, its list of

documents needs to be retrieved and meta data checked. Calculating n_d is performed by taking samples of stored documents and comparing them to other nodes storing the same document.

The basis to make storage and retrieval of meta data secure is to find a path of trusted peers to those peers in the meta data, similar to a web of trust. Three potential attack scenarios are identified: modifying votes, adding votes, and dropping votes. First, it is difficult to modify a vote, since all votes are signed by the respective voter. Second, votes cannot be added or dropped, since a mediator signs them. Thus, if a trusted path exists to the mediator peer and the voter peers, all the above attack scenarios can be detected.

IV. DISCUSSION, CONCLUSION, AND FUTURE WORK

The proposed score-based incentive scheme encourages peers to achieve a high score, as only this allows them to create and approve new documents, as well as to block documents for editing or remove documents in a P2P collaboration system. These tasks are carried out with special modification proposals, which can be either voted on or applied immediately after verifying the score of the submitting peer. The incentive scheme integrates nicely with user-based voting and addresses current issues of PeerVote. Additionally, the score can also be used as a recommendation for undecided voters.

Although the score is published by a peer itself, the score can be verified by any peer. Verification can be exhaustive or sample-based to limit network overhead if many verifications are required. Peers have an incentive to achieve a high score, which allows them to create new documents. However, it is not recommended to use the score as a weight for a vote itself, as this would require each voter to search and retrieve the meta data of a document in order to verify the score, which would generate a lot of additional network traffic, even for the sample-based approach. Future work will focus on the implementation and deployment of the score-based incentive scheme and its evaluation in a real world environment.

ACKNOWLEDGEMENT

This work has been performed partially in the framework of the EU IST Project EC-GIN (FP6-2006-IST-045256) as well as of the EU IST NoE EMANICS (FP6-2004-IST-026854).

REFERENCES

- [1] T. B. Adler, L. de Alfaro: A Content-driven Reputation System for the Wikipedia. ACM WWW 2007, New York, NY, USA, May 2007.
- [2] T. Bocek, B. Stiller. Peer-to-Peer Large-scale Collaborative Storage Networks. ACM AIMS 2007, Oslo, Norway, July 2007.
- [3] T. Bocek, D. Peric, F.V. Hecht, D. Hausheer, B. Stiller. PeerVote: A Decentralized Voting Mechanism for P2P Collaboration Systems, IFIP/ACM AIMS 2009, University of Twente, The Netherlands, July 2009.
- [4] EmanicsLab. <http://www.emanicslab.org/>. March 2009.
- [5] P. M. Johnson. A Glossary of Political Economy Terms. <http://www.auburn.edu/~johnspm/gloss/incentive>. March 2009.
- [6] J. C. Morris, C. Lüer: DistriWiki: A Distributed Peer-to-Peer Wiki. ACM WikiSym 2007, Montreal, Quebec, Canada, October 2007.
- [7] P. Mukherjee, C. Leng, A. Schürr: Piki — A Peer-to-Peer based Wiki Engine. IEEE P2P 2008, Washington, DC, USA, September 2008.