A Reputation Scheme for Peer-to-Peer Media Streaming

Francisco de Asís López-Fuentes[1]

Networks and Distributed Systems Group
[1]Departamento de Tecnologías de la Información
Universidad Autónoma Metropolitana
Av. Constituyentes 1054,
11950 México, D. F., México
{flopez}@correo.cua.uam.mx

Abstract. A challenge in peer-to-peer media streaming systems is how to select good peers in order to realize high quality streaming sessions. The selection of good peers can offer a manner to improve the quality of service via an optimal search or an efficient content delivery. This paper presents an approach for evaluating the participating peers based on their donated resources and on their behavior. This approach uses a reputation/incentives model to isolate the misbehaving peers and the non-cooperating peers, as a way to improve the system performance. Every peer builds its best path using a best-neighbor policy within its neighborhood. The search is based on the best path. A structure based on reputation/incentives policies is used by the supplying peer as a way to assign its outgoing bandwidth to the requesting peers during the media transmission phase.

Keywords: reputation, incentives, peer-to-peer systems, media streaming.

1 Introduction

During the last years, content delivery over the Internet has gained significant popularity. For example, several applications such as TV over IP, streaming and multimedia live streaming require content delivery from one-source to multiple receiver-nodes. On the other hand, peer-to-peer (P2P) networks have attracted the attention from the research community who find in these systems a fast and efficient way to deliver movies, music or software files. A P2P communication infrastructure is formed by a group of nodes located in a physical network. These nodes build a network abstraction on top of the physical network, known as an *overlay network*, which is independent of the underlying physical network with regard to the P2P procedures. An important advantage of P2P systems is that all available resources are provided by the peers. In a P2P system each peer can take the role of both, a server and of a client at the same time. During media distribution, peers contribute their resources to relay the media to others. Thus, as a new peer arrives to the P2P system the demand is increased, but the overall capacity too. This is not possible in a client-server model with a fixed number of servers.

Advances in Computer Science and Electronic Systems
Research in Computing Science 52, 2011, pp. 132-141
Received 23/01/11
Accepted 20/02/11
Final version 07/03/11

A proper selection of peers can lead to the attainment of a good quality of service in terms of a faster search, and a faster distribution of content; but, P2P systems can be affected by misbehaving (or *free-riding*) peers, which reduce the system performance. The *reputation management systems* (RMS) are methods which alleviate this problem [3, 6, 7] through the proposal of appropriate peers, leading, with this, to a natural isolation of misbehaving or non-cooperative peers. A RMS system allows individual peers to rate one to each other according to their past experience with each other. Once a peer has been rated, its rating can be used by other peers to find the best sources of good and authentic content, keeping, at the same time, the effects of malicious peers on the network to a minimum. The proposals for these systems include solutions for the management of trust and the computation or reputation. RM Systems provide a way for building trust without trusted third parties in P2P networks [6].

This paper proposes a method for reputation computation, which involves the concept of incentives. The interest of mixing these two characteristics in this method is the following. First, peers with high reputation can cooperate to make an optimal search or a better content delivery. Second, an incentives-system can encourage the collaboration and exchange of data between peers [1, 5]. Finally, the isolation of misbehaving or non-cooperative peers can avoid the degradation system performance.

The remainder of this paper is organized as follows. Section 2 presents the model and its assumptions. Then, the protocol of the method is presented in Section 3. Section 4 describes the evaluation of our proposed model and presents the results. Section 5 concludes the paper.

2 The Proposed Method

Locating a content-supplier does no guarantee that the service of this supplier will satisfy the user [2], because some misbehaving peers may offer false information in order to maintain a cooperation impression. In order to minimize the effects of misbehaving peers these are to be detected and isolated from the system. Reputation and incentives strategies have been used in several approaches but in separate ways. The concept of reputation is used in several systems such as the online auction system eBay. In eBay's reputation system, buyers and sellers can rate each other after a transaction, and the overall reputation of a participant is the sum of these ratings over the past 6 months. The authors in [3, 9, 10 and 11] are proposing reputation systems with the purpose to ensure that peers obtain reliable information about the quality of the resources they are receiving.

The model here presented makes use of a special peer called manager-peer, which manages the reputation of all peers in the system, and considers that each one of the other peers in the system has information-reputation only of its neighbor-peers. Each peer has local table which keeps a reputation score. Every peer exchanges its local table with any others peers located at no more than 2 hops away. If a peer can be reached, from another one, in just one hop, then it is said that there is a direct link between these peers. If the number of hops is not 1 but 2, then it is said that there is an indirect-link between these peers. The presented reputation method, with incentives, consists of two parts: 1- The reputation model, and 2- The incentives model.

2.1 Reputation Model

The reputation model considers that all peers contribute with their resources to the system. The method uses two components to obtain the average reputation-score in every peer, which are: its capability and its behavior. The first component evaluates the resources of the participating peer: upload capacity, processing capacity, memory, storage capacity and number of shared files. The initial weight for each donated resources can be agreed upon by the users. The second component evaluates the behavior of the participating peer in a cooperation environment, assigning the peer a cheating level a transient level. The cheating level is assigned considering that a peer is cheating when it supplies a wrong content or when it serves with fewer or smaller resources with regard to those promised. The transient level for a peer is determined by the average length of time the peer remains in the system (service-time) and by the average length of time that it takes for the peer to return to the system after it has left. Users could be satisfied when they received content from peers with big resources and good behavior; in the other hand, users could have a bad experience when the involved peers offer low bandwidth, high error-rates, limited processing resources, or frequent disconnections. Every component in the reputation-scheme contributes with its weight, when building the final score.

For example, during a streaming session of a participating peer, the weight of the donated-bandwidth resource of the peer could be bigger than that of its available-storage resource. Initially, the method's reputation score of a peer p_i is based only on its donated resources. As time passes by, if the peer is still connected, it doesn't cheat and its stability and availability is good, then its behavior reputation is increased; otherwise its behavior reputation is to decrease.

The evaluation of the behavior-reputation considers that a transaction made by any peer can be either, performed correctly or not. Our behavior reputation scheme is based on a reputation scheme introduced in [8]. Peers interact using a reputation-approach. A complaint message is evaluated in every peer and in the manager-peer. Any peer, in order to compute the reputation of another peer, evaluates the experience of that peer's neighbors. This is indeed a distributed reputation system, in the neighborhood and in the system. This reputation is usually based on an aggregate of the feedback ratings issued by the diverse peers [6]. The manager-peer is consulted by a peer who wants to know the reputation of peers outside its neighborhood.

The behavior-reputation scheme has two scenarios. First scenario is shown in figure 1. In this scenario, peer p_1 interacts with peer p_2, p_1 may rate the transaction as satisfactory $(t(p_1,p_2)= 1)$ when a retrieval process with p_2 is successful or unsatisfactory $(t(p_2,p_1) = -1)$ if the file is no authentic, if p_2 give false information about its resources or if the download is interrupted. Thus, when peer p_2 is cheating in a retrieval process with peer p_1, a complaint (-1) is sent by p_1 to the global reputation in the manager peer and another complaint is recorded in its local table and distributed on its local neighborhood. A *reputation matrix* is built based on the total number of peers in the neighborhood or in the network. Every peer records the reputation score in the *reputation matrix*, as a local table, while the manager peer records the *reputation matrix* as a global table.

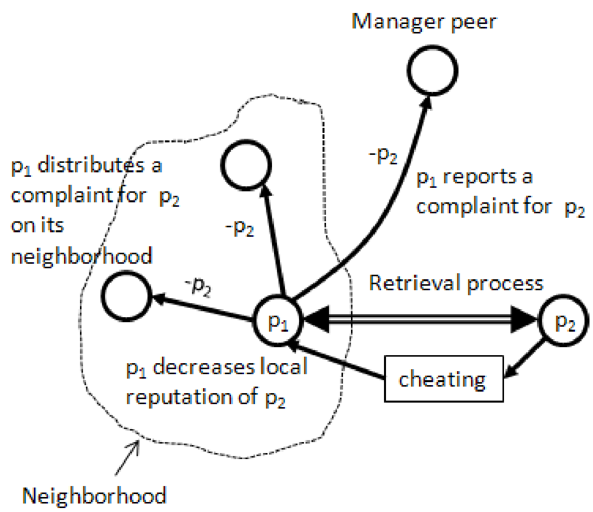

Fig. 1. First scenario for our behavior reputation scheme

Second scenario considers when a new peer joints to the system (see Figure 2). In this case, all entries of new peer are undefined, but these are update as the peer interacts with each other. Every peer updates the reputation of its local reputation matrix, while the reputation of remote peers can be derived from the manager peer.

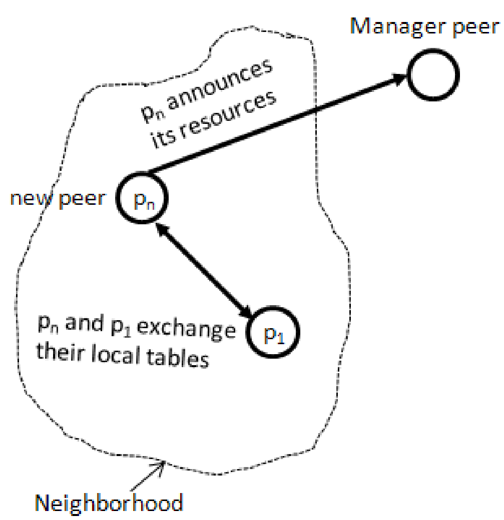

Fig. 2. Second scenario for our behavior reputation scheme

To compute the behavior reputation, the protocol periodically runs an update process in order to update the network. Then, a reputation agent updates the reputation score and the incentives of every peer based on its behavior (cheating level and transient level). Initially, the global reputation score is based on resources only, and the behavior reputation is valued in 0. The behavior reputation score of a peer is increased if it maintains a good service time or it does not cheat, on the other way the behavior reputation score is decreased. In our model the reputation range is between 0 and 10. Where 0 indicates the minimum reputation score and 10 defines the maximum reputation score. We consider the peer status (UP/DOWN/CHEATING) to define several scenarios and to update its reputation score, whenever round expires and the update process start. If the peer status is UP and it does not cheat, the reputation agent computes the average number of rounds that a peer remains connected to the network and its behavior reputation score. If the peer status is DOWN and it does not cheat, we need to determine how long it is DOWN. If time is greater than 4 rounds, then the peer will be punished in 2 rounds when it returns to the P2P network (UP status). During the first round its reputation score will be zero and in the second round its reputation score will be based on resources only. When the peer status is CHEATING, its reputation is decreased to 0 in all its neighbor peers and in the manager peer. All peers isolate the cheating node, and they do not send, forward or receive any messages or packets from it.

2.2 Incentive Model

An incentives model needs to define some rules to motivate peers to contributes more resources and avoid the no-cooperation in the system [12, 13]. For our incentives system we assume the following rules:

- A peer cooperates with another peer based on its generosity factor [4].
- Peers with high reputation are allocated close to the source, forming rings. The internal ring will be close to the source and it will have the highest reputation.
- Peers with high reputation receive a high priority to upload contents during a contention.

We adopt an approach based on the game theory to address the no-cooperation problem in the system. In particular, we use a choking algorithms model to capture the essential tension between individual and social utility, asymmetric payoff matrices to allow asymmetric transactions between peers, and a learning-based population dynamic model to specify the behavior of individual peers, which can be changed continuously. So, our approach rewards cooperation and therefore considers upload and download rate, like a generosity factor, to translate cooperation to earn benefits or loss it to peers which cooperate or not.

Our approach considers the upload and download rate and rewards the cooperation using a generosity factor. The peer's cooperation is translated to benefits via the generosity factor. This factor measures the benefit that a peer provided relative to the benefit it obtained. On this way we try to avoid that the system collapses when we have peers consuming more services than they provide [4]. The generosity resumes the General Prisoner's Dilemma for an asymmetric payoff matrix [4, 5]. Let Up_i and Dp_i the provided service and the obtained services by the cooperating peer p_i, respectively. A unit of provided service is a packet given successful, while an unit of obtained service is a packet received successful. Then, the p_i's generosity is given by

$$G(p_i) = Up_i / Dp_i \tag{1}$$

To reach an effective cooperation in the neighborhood, every peer uses its own generosity as a measuring stick to judge its peer's normalized generosity [4] expressed by:

$$G_{p_j}(p_i) = G(p_i) / G(p_j) \tag{2}$$

where $G(p_j)$ is the p_j's generosity and $Gp_j(p_i)$ measures p_i's generosity relative to p_j's generosity.

Using its reputation table each peer builds a hierarchical structure based on rings. We call them, reputation rings. The reputation rings is used by a supplying peer during the download phase. Peers with high reputation are allocated close to the supplying peer, while the peers with low reputation are allocated in the extern rings and distant to the source. A requesting peer with high reputation receives a high access priority (download) from the supplying peer during a contention. The number of reputation rings and its reputation thresholds are values that the system architect is free to set. A requesting peer p_i is allocated in a reputation ring R_i based in its reputation score. The incentives are distributed among peers based on this hierarchical structure. Peers close to the source receive more incentives than peers far from it. In our approach, the bandwidth is considered the main incentive. We distributed a different source bandwidth portion among the reputation ring. Thus, peers allocated into the internal reputation rings receive a bigger bandwidth portion from the supplying peer than peers allocated into the external reputation rings. We give a different percentage of incentives to every reputation ring. These percentages can be arranged by users of the system. With the reputation and the generosity factor, each source peer encloses its requesting peers in reputation/generosity rings.

Our proposed model updates the reputation of all enrolled peers running periodically a reputation process, which can be activated by every participating peer or by the manager. This process collects information about the updated resources and behavior of each peer within its neighborhood, and updates the local table. Our protocol operates in two phases, the first is called search phase and the second is the download or streaming phase.

3 Evaluation

We evaluate our reputation mechanism with incentives using a simulator based on java. A random topology is generated and different traffic scenarios are used. Thus, the simulation considers two mainly files. On the one hand the topology, and on the other hand, the traffic scenario file. Our experiments use a random topology with 40 nodes. Additionally, the resources donated by each peer such as CPU, bandwidth, memory and storage capacity are recorded. The initial reputation is based on resources only. Also, a manager peer is defined in the topology.

Regarding to the traffic scenario, we develop a generator of traffic that generate a traffic scenario based on the real internet traffic. We assume that the peer arrives and leaves (UP/DOWN) to the system following a Poisson distribution. Also, a Poisson distribution is used to model the time that a peer remains in the system. Query popularity (number of queries and time between queries) is modeled following a Zipf distribution. Every peer selects a set of videos in a random uniform way. In the beginning, there are no relationships among peers in the system. The simulator reads the topology and the traffic scenario files and traffic events such as UP's, DOWN's, CHEATING's, Queries are generated during the simulation. Every node initializes its local table with information about its neighbors and defines its shared files. The reputation mechanism collects the donated resources by each peer and monitors its behavior while the incentives system distributes the incentives based on the rules described in section 2.2. To simulate the incentives over the reputation rings, we established our hierarchical structure in 5 rings, and the reputation is distributed among them. We fix a bandwidth portion of 30%, 25%, 20%, 15% and 10% for the rings 1, 2, 3, 4 and 5, respectively. In our simulation, the peer's bandwidth is considered the main resource and its changing behavior is considered in both phases. Other resources such as storage, CPU, memory and number of files retain its initial values in the system.

The results show how the system performance is greatly affected when the rate of the misbehaving peer is increased. We measure the system performance as the number of successfully downloaded contents. We simulate the system performance using different percentage of misbehaving peer such as 0%, 15%, 30%, 45%, 60% and 75%. Figure 3 shows how the system performance is affected by the misbehaving peers. Here, we can see that the number of download content is decreased proportionally to the number of misbehaving peers. We use two alternatives, in the first alternative we sent a query using a best path only, and the second alternative each peer allocated within the query route uses its two best paths.

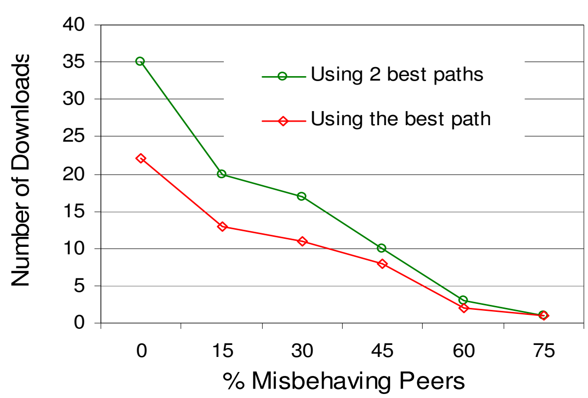

Fig. 3. System performance is affected by the misbehaving peers

Contrary, Figure 4 depicts the system performance when the cooperation between peers is increased. We can see that more files are fully downloaded when the cooperating peer percentage is increased.

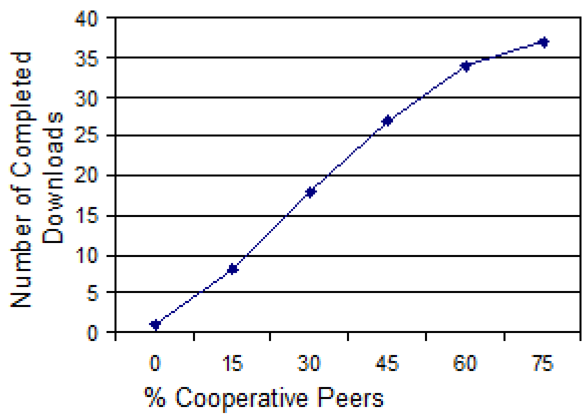

Fig. 4. Cooperation between peers increases the system performance.

Finally, our mechanism isolates misbehaving peers from good peers. On this way, any misbehaving peer cannot process, send or forward any content or query. Either the good peers cannot send queries to the misbehaving peers or receive requests from them. Figure 5 shows how a system without reputation allows download content from the misbehaving peers increasing the probability of having a greater number of corrupted contents. In this scenario we can see that misbehaving peers give bad ratings and good peers always give good ratings in the system. However, in the real electronic communities correcting the malicious peer's behavior is a hard task. Instead of correcting each such malicious peer, we need to minimize its impact in the system performance.

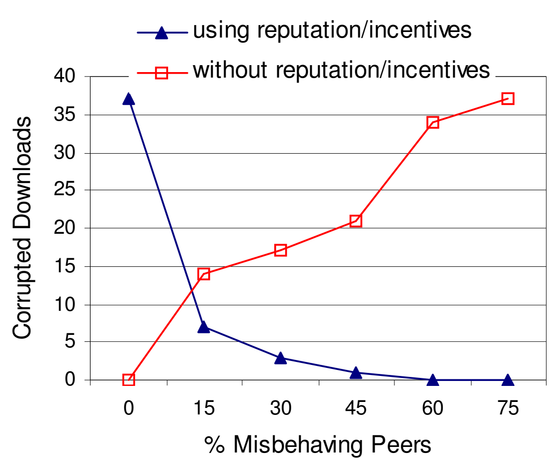

Fig. 5. Comparison of rejected corrupted files

4 Conclusions

In this contribution, we have proposed and evaluated a reputation mechanism with incentives for a P2P system. Most of the reputation systems consider correction of malicious peers by giving incentives for positive feedbacks. However, in our proposed model isolates misbehaving peers from good peers, and incentives are only used to gain most cooperation in the system. We show how the presence of misbehaving peers reduces the system performance. Our results show that using our proposed mechanism, free riding can be reduced, because the non-cooperating peers are eliminated from the system. Also, the reception of corrupted files from the misbehaving peers is eliminated. Using the reputation rings a peer can distribute its upload capacity among good peers based on its reputation score. The possible extensions for this work could be in the direction of large P2P streaming systems, where fully distributed and scalable schemes must be improved. Finally, we will perform more evaluations in order to compare our proposed scheme with other methods available in the literature.

References

1. Silverston, T., Fourmaux, O., Crowcroft, J.: Towards an Incentive Mechanism for Peer-to-Peer Multimedia Live Streaming Systems. In: Th IEEE Eight International Conference on Peer-to-Peer Computing, pp. 125--128, 2008, IEEE Press, New York (2008).
2. Marti, S., Garcia-Molina, H.: Limited Reputation Sharing in P2P Systems. In: 5th ACM Conference on Electronic Commerce, pp. 91--101. ACM press, New York (2004).
3. Kamvar, S. D., Schlosser, M. T., García-Molina, H.: The EigenTrust Algorithm for Reputation Management in P2P Networks. In: 12th International World Wide Web Conference, pp. 20-24, ACM press, New York (2003).
4. Feldman, M., Lai, K., Chuang, J., Stoica, I.: Robust Incentive Technique for Peer-to-Peer Networks. In: 5th ACM Conference on Electronic Commerce, pp. 102--111, ACM Press, New York (2004).
5. Blanc, A., Liu, Y. K., Vahdat, A.: Designing Incentives for Peer-to-Peer Routing. In: 25th IEEE International Conference on Computer Communications, pp. 374--385, vol. 1, IEEE Press, New York (2005).
6. Prasad RVVSV, Srinivas, V., Kumari, V. V., Raju, KVSVN: An Effective Calculation of Reputation in P2P Networks. In: Journal of Networks, Vol. 4, No. 5, pp. 332--342, Academic Publisher, West Bengal, India (2009).
7. Su, X., Dhaliwal S. K.: Incentives Mechanisms in P2P Media Streaming Systems. In: IEEE Internet Computing, pp. 74--81, IEEE Press, New York (2010).
8. López-Fuentes F. A. and Steinbach E.; Architecture for Media Streaming Delivery over P2P Networks. In: Springer-Verlag, LNCS 3563, pp. 72-82, Berlin Heidelberg, Germany (2005).
9. Cornelli, F., Damiani, E., Vimercati, S.D.C.D., Paraboschi, S., Samarat, S.: Choosing Reputable Servents in a P2P Network. In: Proceedings of the 11th World Wide Web Conference, Hawaii, USA (2002).
10. Gupta, M., Judge, P., Ammar, M.: A Reputation System for Peer-to-Peer Systems. In: Proceedings of ACM NOSSDAV 2003, Monterey, CA (2003).

11. Aberer, K., Despotovic, Z.: Managing Trust in a Peer-2-Peer Information System. In: Proceedings of the ACM 10th International Conference on Information and Knowledge Management, New York, USA (2001).
12. Habib A., Chuang, J.: Incentive Mechanism for Peer-to-Peer Media Streaming. In: Proceedings of the 12th IEEE International Workshop on Quality of Service, Montreal, Canada, IEEE Press, New York (2004).
13. Lai K., Feldman M., Stoica I., Chuang J.: Incentives for Cooperation in Peer-to-Peer Networks. In: Proceedings of the Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA (2003).