J. Parallel Distrib. Comput. 72 (2012) 809–818

Contents lists available at SciVerse ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

The importance of considering unauthentic transactions in trust management systems

Pilar Manzanares-Lopez *, Josemaria Malgosa-Sanahuja, Juan Pedro Muñoz-Gea

Department of Information Technologies and Communications, Technical University of Cartagena, Antiguo Cuartel de Antigones Campus Muralla del Mar s/n, E-30202 Cartagena, Spain

ARTICLE INFO

Article history:
Received 24 May 2011
Received in revised form
5 March 2012
Accepted 13 March 2012
Available online 23 March 2012

Keywords:
Trust management
P2P systems
File sharing

ABSTRACT

In peer-to-peer (P2P) networks, trust management is a key tool to minimize the impact of malicious nodes. EigenTrust is claimed to be one of the most powerful distributed reputation management systems focused on P2P file-sharing applications. It is the theoretical base of other systems, and it has also been directly modified in an attempt to improve its performance. However, none of them give appropriate importance to all the information about transactions. This paper proposes an enhancement of EigenTrust, which considers unsatisfactory transactions in greater depth. Pos&Neg EigenTrust is able to obtain a blacklist of the identities of the malicious nodes. Therefore, it is able to significantly reduce the number of unsatisfactory transactions in the network.

1. Introduction

Two types of adversary can be identified in peer-to-peer (P2P) networks: selfish peers and malicious peers. Whereas the former (known as *free riders* in file-sharing networks [10]) use system services contributing minimal resources or nothing by themselves, the latter cause harm to either specific targeted members of the network or the system as a whole (for example, distributing corrupted or unauthentic files in file-sharing networks).

In today's P2P networks, incentive schemes [5,8,15] can be used to encourage cooperation of selfish peers, but they can be ineffective against malicious peers. Trust management systems are a key tool to minimize the impact of malicious node actions [13]. Here, trust is understood as the confidence that one peer has about another peer's reputation. The interactions among peers give the opportunity to acquire the necessary experience (measured in terms of good and bad transactions) to build up a reputation management algorithm. Basically, these algorithms are able to calculate and share among the network community a global trust value for each peer.

Several representative P2P reputation systems have been proposed in recent years. Some of them are focused on trust management in P2P e-commerce applications, while others are focused on generic P2P applications such as P2P file sharing. Belonging to the latter, EigenTrust [9] is claimed to be one of the most powerful reputation management algorithms. This is, in greater or lesser extent, the theoretical base of many other P2P reputation models.

EigenTrust is a distributed protocol that is able to compute the global trust value of a peer by applying each peer's local trust values transitively. To guarantee the convergence of the global trust values, EigenTrust uses linear algebraic theory. This system is focused on minimizing the number of unauthentic downloads during the transactions. To reach this aim, EigenTrust focus its work on the identification of the peers that offer more authentic files. A perfect balance between simplicity and efficiency has made EigenTrust the most widely used reputation management algorithm up to now. In fact, the vast acceptance of this system in the research community is patently obvious in the set of proposals trying to improve its performances.

Although identifying a peer with good behavior is a key task to reduce the number of unauthentic downloads, in a real scenario in which different threats can be caused by other peers, it can be as interesting, or indeed more interesting, also to identify these malicious peers. They introduce unauthentic files in the system and also introduce wrong information about the transactions with others, adding additional noise to the system.

To our knowledge, none of the EigenTrust-based systems deal with this aim. All of them present a limitation as regards the definition of local trust values in the original EigenTrust: it does not take into account the cases where the difference between satisfactory and unsatisfactory transactions is negative. These negative values can arise in two ways: peers can upload (whether consciously or not consciously) unauthentic files, and also peers can lie about the evaluation of the transactions. If the negative

* Corresponding author.
E-mail addresses: pilar.manzanares@upct.es (P. Manzanares-Lopez), josem.malgosa@upct.es (J. Malgosa-Sanahuja), juanp.gea@upct.es (J.P. Muñoz-Gea).

0743-7315/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
doi:10.1016/j.jpdc.2012.03.006

values were considered, one of the most important properties of the algebraic model would not be met. Therefore, it would be impossible to guarantee the convergence of the global trust values. In other words, EigenTrust-based systems are able to measure the goodness of the peers but they cannot measure their wickedness. This behavior can be exploited by the malicious peers to improve their efficiency in shutting down the system.

Our work proposes an enhancement of EigenTrust called Pos&Neg EigenTrust. This system considers all the information about transactions in more depth. However, it is defined so that all the algebraic properties of the original are still valid. After some transactions, Pos&Neg EigenTrust can obtain two ordered lists. The first list classifies the peers on the basis of their goodness and the other list classifies them on the basis of their wickedness. Using these rankings adequately, it is possible to build a blacklist containing the identities of the malicious nodes. This is a significant improvement over the original EigenTrust, which is focused only on the identification of good peers.

This paper evaluates our proposal by simulation of two types of threat model. The first one (a simple threat model called individual malicious peers) consists of completely isolated nodes which always provide unauthentic files. The second one (a more complex model called malicious peers in a collective with camouflage) is more realistic. The malicious peers work collectively and sometimes (maybe often, maybe rarely) provide unauthentic files (that is, they are camouflaged). When evaluating that last type of malicious node, two attacks are considered. In the first case, malicious peers always lie about the local trust values associated with the rest of the peers. In the second one, a more sophisticated attack, the malicious peers only lie with a certain probability. Simulation results show that, thanks to the identification of good nodes and also the identification of malicious ones, Pos&Neg EigenTrust is able to significantly reduce the number of unsatisfactory transactions in the network with any type of malicious nodes.

The rest of the paper is organized as follows. Section 2 describes previous works on P2P reputation management systems, besides identifying the contribution of this paper. Section 3 summarizes the original EigenTrust system and describes its main shortcomings. Section 4 presents our proposal. In Section 5, both the EigenTrust and the Pos&Neg EigenTrust systems are evaluated and compared by simulation. Section 6 concludes the paper.

2. Related works in P2P reputation systems

Trust and reputation models have been proposed as a solution for guaranteeing a minimum level of security between entities belonging to a distributed system (such as P2P networks, ad hoc networks, wireless sensor networks, or multi-agent systems) that want to have a transaction or interact [6].

In the case of P2P networks, in which users do not know each other at all or, at least, do not know everyone, the application of trust mechanisms can help peers to find out which is the most trustworthy peer to interact with, thereby preventing or minimizing the selection of a malicious one.

Most of the current P2P reputation systems follow the same general steps. Peers obtain information about others in the community from their own experience and/or asking other users their opinions about the others. Then, the obtained information is aggregated properly and somehow a score for every peer in the network is computed. Thus, the most trustworthy or reputable entity in the system providing the required service is selected to have an interaction with it, evaluating a posteriori the satisfaction of the received service. Finally, according to the satisfaction obtained, the trust or the reputation deposited in the selected service peer is consequently adjusted. How these general steps are implemented and how the trust parameters are defined differ from one reputation system to another.

As mentioned before, the classification of P2P reputation system depends on the applications they are focused on. Next, this section summarizes some of the most representative reputation models for file-sharing P2P systems.

EigenTrust [9], described in detail in the next section, is one of the most representative global reputation P2P systems. It is one of the most known and cited reputation models besides being, to a greater or lesser extent, the theoretical base of other models. It relies on linear algebraic theory to guarantee obtaining global trust values. Moreover, it is very simple and does not depend on complex parameters or assumptions. For these reasons, it is one of the most widely studied trust management systems.

EigenTrust assumes that all the peers perform transactions with others evenly. However, in practice, there are some closed community groups in which the global trust values used by EigenTrust do not adequately reflect the trustworthiness of peers [11]. In these scenarios, the EigenTrust reputation model will not be able to identify the malicious peers and, consequently, will not be able to isolate them from the network (for an interesting survey of security threats scenarios in P2P reputation models, see [7]).

It is widely accepted that, in order to solve this problem, it is necessary to calculate the global trust values taking into account the fact that peers can be classified into groups or profiles (at least two profiles: good and malicious). For example, [12,17] propose calculating the global trust values weighting the local trust values by a parameter called distance (Euclidean and Hamming distance, respectively), which in turn, is a measure of the similarity among different profiles. On the other hand, [2,1], and, more specifically, [18] propose the use of more complex local aggregation functions which consider the existence of different groups. [18] demonstrates that, in a general P2P reputation system, the number of feedbacks follows a power-law distribution (only a handful of peers have the most feedback) [4]. This fact enables the definition of a complex local aggregation algorithm with an overhead which grows only linearly. In addition, this handful of peers is dynamically selected to carry out special tasks to reinforce the reputation network.

However, the above problem can also be faced from a different point of view. Neither EigenTrust nor any of its improvements consider the limitation that, in our opinion, has the local trust value definition. As analyzed in detail in Section 3.4, EigenTrust, and also [12,17,2], make limited use of the information about unsatisfactory transactions. They only consider the cases where the difference between satisfactory and unsatisfactory transactions is positive. Other works such as [18] do not impose a specific algorithm to calculate local trust values, enabling each peer to have its own criteria to generate feedback scores. Other cases, such as [1], do not even consider unauthentic transactions to calculate local trust values.

Our proposal improves the original EigenTrust system by considering all the transaction information in more detail. Pos&Neg EigenTrust uses both positive and negative values of the difference between satisfactory and unsatisfactory transactions when obtaining the local trust values. Thus, in addition to the global trust vector, another useful vector called the *negative trust vector* or *bad reputation vector* is obtained. Using both adequately, it is possible to obtain a blacklist containing the identities of the malicious nodes, whatever their behavior is.

To our knowledge, this approach has only been taken into account in [3]. There, the authors define new metrics that aggregate the negative opinions expressed by peers: badness, positive dishonesty, and negative dishonesty. They define two methods to integrate these metrics into EigenTrust: the integration