Reputation management in distributed systems

Article

Accepted version

A. De Paola, A. Tamburo

In Proceedings of the 3rd International Symposium on Communications, Control and Signal Processing, 2008. ISCCSP 2008, pp. 666-670

Publisher: IEEE

http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4537308

Reputation Management in Distributed Systems

Alessandra De Paola and Adriano Tamburo
Dipartimento di Ingegneria Informatica,
Università degli Studi di Palermo
Viale delle Scienze, 90128 Palermo, Italy
e-mail: {depaola, tamburo}@csai.unipa.it

***Abstract*—Several distributed applications, implemented over today's Internet, are based on the assumption that participating agents collaborate in order to achieve their own goal. However, when these applications are modelled as unstructured distributed systems, the greater autonomy and decentralization encourage antisocial behaviours, which are likely to cause performance degradation for the whole system. This paper presents a fully distributed reputation management system that allows the evaluation of agent reputation in unstructured environments without any centralized coordination. The proposed approach is based on game theory and is capable of capturing the highly dynamic nature of the involved communities. As a representative example of an unstructured environment, Peer-to-peer (P2P) networks are considered. Those dynamic communities are affected by several antisocial behaviours, such as free riding. Since this phenomenon typically causes and exacerbates an unbalanced and unfair use of system resources, it has been considered as the case study in our work. The proposed solution exploits peer reputations in order to define an incentive system, whose main goal is the dissuasion from free riding.**

I. INTRODUCTION

In distributed environments where heterogeneous agents collaborate and interact in order to achieve their own goals, obtaining guarantees on their behaviour from some reliable central authority is typically unfeasible. In such scenarios, a reputation management system capable of building a profile representing the reliability of each agent may prove extremely useful. Knowing agent reputation is particularly helpful for detecting those agents that are deceitful or potentially dangerous for the community; however, reputation management becomes a very challenging task, when the distributed system is fully unstructured. In this case, interactions among agents are unpredictable, the community is often highly dynamic, and no central authority is present to carry out supervision and coordination activities. The lack of a supervising entity favours the arousal of antisocial behaviours: several agents may act opportunistically, disregarding community cooperativeness. Selfish behaviour cannot be regarded as an anomaly, rather it can be explained from the point of view of particular convenience.

According to game theory, the analysis of agent interactions cannot take "altruism" into account. Each agent selects its own actions in order to obtain the greater advantage, to the best of its knowledge, even to the detriment of the whole community. In order to contrast antisocial behaviours, some interventions are required to modify the utility as perceived by agents, so that not collaborating may be neglected in favour of cooperation. Game theory, and in particular its branch called evolutionary game theory [1], provides appropriate tools for modelling such complex systems and for verifying formal properties of incentive mechanisms specifically designed for contrasting antisocial behaviours.

A perspicuous example of distributed systems where such selfish behaviours may have a negative impact on the overall performance is represented by P2P networks; in those distributed systems, agents with equivalent functionalities interact without any centralized coordination. Unstructured P2P networks, in particular, achieve the highest decentralization degree: the network is self-organized without any predefined overlay structure. Unfortunately, the lack of a centralized control favours the occurrence of free riding phenomena, characterized by a resource exploitation by a few system agents; for instance, in file-sharing applications, free riders share few or scarcely appealing resources. Several recent works [2], [3] have analyzed the relevance of this phenomenon in different models of P2P networks such, for instance, the Gnutella network. The general conclusion is its potential threat for the survival of the network itself. It modifies the network structure, that degenerates into a client/server architecture, with the consequent performances decrease. Free riders deliberately do not cooperate since collaboration implies some costs, and this behaviour reduces the perceived utility. In P2P networks costs are related to the resources that a peer has to make available in order to actively connect into the system; among these, link bandwidth, CPU time of the hosts where P2P daemons run, and storage space. Such resources are used both when a peer directly answers in order to provide the requested resource, and also when a peer simply collaborates in order to find some resources, routing queries coming from other peers.

This paper proposes a distributed reputation management scheme, where participating agents interact in order to build a view of the network that is as homogeneous as possible with respect to agents reputation. The reputation management mechanism begins with a local assessment and then, through a reputation diffusion algorithm, propagates information among cluster of mutual reliable agents. The proposed system has been tested on an unstructured P2P network that adopts a flooding-based routing protocol for resource localization.

The remainder of the paper is organized as follows: in Section II other works presented in literature are described, both related to the general issue of reputation management in distributed systems, and to the particular problem of contrasting the free riding phenomenon in P2P networks; in Section III some basic principles of game theory are summarized and the choice of the evolutionary game theory, as the best choice for these complex systems, is motivated; Section IV describes the proposed reputation management system and Section V reports the experimental results obtained by preliminary simulations; finally, Section VI states some conclusions and outlines possible future developments of this work.

II. Related Works

A. Reputation Management in Distributed Environments

In the past few years, several solutions have been proposed in order to provide reputation management in distributed environments. Most of the proposed approaches aim to evaluate agent reliability in a generic distributed system. Reputation management systems can be classified into *centralized* and *distributed* approaches.

In the *centralized* approach, a central reliable entity monitors transactions among agents in the community and computes a reputation value for each of them. Typical applications of centralized reputation systems are e-commerce systems, such as eBay, in which after any transaction, the buyer assigns a score in order to evaluate seller reliability. Epinions [4] introduces the concept of double reputation: a reputation value for each agent providing a service, and a reputation value for each user expressing an opinion about a service. The idea is very interesting, since users with higher reliability have more influence during the overall reliability evaluation of a service or product; this way the reputation value is highly reliable. However, both approaches may prove unpractical in a fully distributed environment.

In the *distributed* approach no centralized control structure is present, rather each agent needs to collect data in order to calculate its own estimation about the reputation of all other agents, which it gets in contact with. This information could be shared among all the agents belonging to the same communities, using several information diffusion mechanisms. The EigenTrust project [5] adopts a distributed approach in order to calculate a global reputation value in a P2P network. It uses the satisfaction level expressed by peers, at the end of each interaction; this value constitutes a local reputation value. Local values are merged in order to provide a single, global reliability parameter. Each peer calculates a reputation value for its neighbour peers, averaging the reputation values provided by peers directly connected. Weights used in the average computation are the reputation values of evaluating peers. The global reputation value is computed by diffusing this evaluation to peers located more hops away from the evaluating peer, again averaging values obtained by information propagation. The proposed system does not explicitly take into account bad reputation: an unknown peer is simply considered a malicious peer.

In [6] a reputation management system based on fuzzy logic is proposed. It describes a system capable of successfully managing uncertain and partial information reported by other network peers. As in previous works, reputation values are at first collected locally and then propagated over the network; the information gathering mechanism is again based on the weighted average of reputation values, although enriched with a fuzzy logic system. Both approaches adopt Distributed Hash Tables (DHT) in order to maintain a unique global reputation. Using a DHT structure involves that a single peer is responsible for holding the reputation value of a given peer; every time this information is requested, the holding peer must be contacted; furthermore, each node failure on the network causes the reallocation of the managed identifiers. Previous considerations make it clear that DHT-based solutions are unpractical for highly dynamic systems such as P2P unstructured networks. In [7], in addiction to the good reputation value, the evaluation of a bad reputation value is proposed. This is a distinguishing feature in comparison to other works proposed in literature. However, authors assume the full collaboration of peers during the information diffusion, and this assumption appears rather unrealistic.

B. Reputation Management for Countering Free Riding

Most of the approaches proposed in literature aim to address free riding by using reputation management systems. These approaches try to estimate the goodness of a peer, as its cooperation degree. After estimating a peer's goodness, it is possible to employ a service differentiation, allowing more collaborative peers to obtain higher quality services. In [8], the authors propose two different incentive systems based on reputation, in which a reputation score is assigned to each user. In the first case, each user is authorized to obtain resources only from users with lower or equal score; in the second one, the reputation index is used in order to give higher priority to requests coming from collaborative users. Both methods require a secure and reliable mechanism to maintain user reputation values. However, the overhead of maintaining and broadcasting this information becomes prohibitive for fully distributed networks. The authors of [9] propose a distributed reputation system in which every peer maintains a reputation value about other peers which it gets in contact with; in this case, peer reputation depends on the provided upload bandwidth. To the best of our knowledge, this is one of the few works where some ideas borrowed from evolutionary game theory are adopted in the context of P2P networks. The total upload bandwidth shared by a peer is estimated through a mechanism based on Simulated Annealing algorithm, and the considered approach tries to estimate the reward associated to each upload bandwidth value.

Our paper relies on the reputation management system proposed in [9], improved and modified in order to manage bad reputation in addition to good reputation, and to allow reputation diffusion.

III. Game Theory Background

Game theory aims to model systems composed of several agents interacting with each other while still trying to achieve their own particular goal, which typically consists in the maximization of a utility function representing the benefits received by each agent in a given situation. A situation is the result of all the actions accomplished by all other agents (called players) involved in the interaction (called game). Therefore, game theory is used to represent scenarios in which an agent cannot completely control the utility it perceives, which also depends on the choices made by other agents.

A particular category of games is represented by non-cooperative ones; in such games, an inefficient equilibrium is reached as a result of all agents making rational choices, i.e. selecting actions maximizing their own utilities. An equilibrium point is inefficient if there exists a different situation in which the utility perceived by each agent is higher. In non-cooperative games, the efficient equilibrium point cannot be reached; namely, in the absence of a full cooperation among players, each agent perceives that not cooperating is more convenient than doing it. In such games no reliable third party assumes the role of forcing all participating agents to cooperate. Unfortunately, this is the case in which most distributed systems fall.

A. Evolutionary Game Theory

Classic game theory studies strategy games, in which agents are considered perfectly rational: each agent knows exactly all actions that other agents are allowed to perform, as well as the utility value perceived by each of them in each situation. A perfectly rational agent is thus able to select the best action that maximizes its own utility, given the choice performed by other agents. Most real-life complex systems violate these theoretical conditions, so classical game theory results unsuitable. In particular, in several scenarios, agents do not know their own utility function, and least of all, the utility function of other agents. In P2P networks, for example, peer utility can be defined as the benefit perceived when obtaining resources from other peers, net of costs requested for cooperation. It is not convenient to express this function in a closed form, too heavily bound to system design choices, and the most appropriate approach consists in learning it during the continuous interaction with other agents.

In distributed environments, in which a wide set of not perfectly rational players repeatedly interact, classical game theory does not appear as the best tool for modelling the system. The most promising tool for modelling these complex systems appears to be a particular branch of game theory, known as evolutionary game theory [10]. It analyzes the behaviour of a population of players. The game is played several times between pairs of players, selected randomly. At the end of each match, the loser is able to change its penalizing strategy, allowing the evolution of the community strategy toward the most rewarding one. This approach allows considering more complex realms, in which agents are not perfectly rational, either because they do not have a complete and exact vision of the system, or even because they do not retain the computing capabilities required to find the best strategy among those possible. In evolutionary games, the best strategy is not typically computed as a closed-form expression; on the contrary, the best strategy emerges from a trial-and-error learning mechanism, thanks to which players can understand that, in a given situation, a given strategy is preferable to another one.

The way players modify their strategy, according to the results of previous matches, is called *evolutionary dynamic*, and several flavours may be considered, among which the most popular is the replicator dynamic. This assumes that the diffusion speed of a given strategy depends directly from the difference between the utility of this strategy and the population average utility. However, this approach does not focus on the learning mechanism of the single player that, as previously seen, does not possess the system global view. For this reason, it results suitable only for an external system analysis. The simplest replicator dynamic, among those focused on the individual choice method, is the imitation dynamic that consists of the trivial imitation of the winning strategy, after each match. A more robust approach may involve the exploitation of the past experience and of the current conditions in order to learn the best strategy. In this context, any learning method can be adopted, such, for instance, the reinforcement learning algorithm.

IV. The Proposed Reputation Management System

This paper proposes a fully distributed reputation management system, in which autonomous agents interact in order to build a reputation value of all participating agents. The aim is to build a reputation estimate that will result as homogenous as possible among all agents. As already mentioned, we adopted as case study the scenario offered by an unstructured P2P network. This class of networks represents a fitting example of distributed environments that are plagued by antisocial behaviours. Furthermore, because of their strong autonomy and dynamicity, P2P networks are not suitable for the adoption of a structured approach, such as the systems based on DHT. The estimated reputation values are used within an incentive system whose aim is the penalization of not cooperating peers. The main feature of our system is its capability to adapt the utility perceived function during the interaction among agents. Through this modification, the cooperating behaviour becomes a more appealing choice. In our system, this means that peers are motivated both to share more precious resources and to actively attending to the query forwarding process. In fact, the greater the cooperation level of a peer, the greater its reputation value maintained by agents with which it has interacted.

The incentive mechanism privileges peers with higher reputation, supplying them with more resources with respect to other system agents. Peers that are considered free riders are penalized, rejecting their requests with a probability proportional to their bad reputation.

In the reputation management mechanism two phases can be distinguished: a former for building a local reputation estimation and a latter where the reputation information is broadcasted in all the community. The first phase (local estimation) exploits the information colleted during the interaction among peers, integrating past values and information about the current context. During the reputation diffusion phase, pairs of mutually reliable peers exchange their local reputation values, so each peer can integrate the received information with its own local reputation value.

A. Local Reputation Estimation

The sub-system in charge for the local reputation estimation, requires that each peer estimates a cooperation degree of other peers with which it has interacted. The purpose of our incentive system is to achieve a stable equilibrium between resources that are provided and requested by an interacting peer. In particular, the proposed mechanism aims to maintain a balance among each pair of peers in the network. The local reputation mechanism uses the difference between supplied and obtained resources from a peer to another one, advantaging both peers in a balanced situation or the more generous peer otherwise. On the other hand, those peers, for which the number of obtained resources results greater than supplied ones, are penalized.

Two different values are used to represent trust and distrust about a peer; in the following these values are called *good reputation* and *bad reputation*. The design choice of adopting the different values was made according to what proposed in [7]. In several works present in the literature, a single reputation value is adopted. However, using only this unique value, often, it may result hard to distinguish between new unknown peers and peers with malicious behaviours. To overcome this difficulty, in our system both the values are considered into a reinforcement-learning scheme that uses the difference between requested and supplied services.

Furthermore, most of the previous systems consider, as supplied services, only successfully completed downloads. The query forwarding mechanism, which indeed constitute the fundamental core of any file sharing P2P application, is not considered. At the best of our knowledge, paper [11] is the only work considering the query routing as a supplied service, although a system based on token-exchange, rather than a reputation-based approach, is adopted. However, using tokens as a virtual coin requires a centralized trust authority. As extensively noted this kind of system is not suitable for unstructured P2P networks.

In our reputation mechanism, a peer maintains a reputation profile for each other peer, which has been previously contacted. The reputation profile maintained by peer i about peer j, contains the number of requests that have been satisfied for the peer j (req_i^j) and the number of services obtained from peer j ($serv_i^j$). More precisely, req_i^j represents the number of responses generated and the number of the query routed, by the evaluating peer for the peer j; $serv_i^j$ is the number of resources obtained and responses routed, by the peer j for the evaluating peer. The above values are used to calculate the local values for the instantaneous estimation of bad and good reputation for peer j ($\hat{gr}_i^j$ e $\hat{br}_i^j$), together with actual reputation values collected so far (gr_i^j e br_i^j).

$$\hat{gr}_i^j = \begin{cases} max(1-\frac{req_i^j}{serv_i^j},\ gr_i^j), & \text{if } serv_i^j \geq req_i^j \\ 0, & \text{otherwise} \end{cases};$$
$$\hat{br}_i^j = \begin{cases} max(1-\frac{serv_i^j}{req_i^j},\ br_i^j), & \text{if } serv_i^j < req_i^j \\ 0, & \text{otherwise} \end{cases}. \quad (1)$$

The actual reputation values are updated on the basis of these instantaneous reputation values, through a reinforcement learning approach. Indicating with $r_i^j(t)$ the generic reputation value at time t (good or bad) held by peer i about peer j, the update is carried out according the following equation:

$$r_i^j(t) = \alpha * \hat{r}_i^j(t) + (1-\alpha) * r_i^j(t-1). \quad (2)$$

B. Reputation Diffusion

The protocol for the reputation diffusion aims to broadcast information about the peer behaviour, among neighbours considered mutually reliable, so as to obtain a view of the network that is as homogeneous as possible, with respect to the reputation values. The diffusion mechanism to some extent mirrors human social interactions: when two agents consider themselves mutually reliable, they exchange information about other members of their community, in order to integrate the directly acquired information. Leveraging only the information coming from reliable agents avoids that malicious agents may attack the reputation system providing forged values. Furthermore, providing the reputation information only to reliable agents constitutes an additional obstacle for non-cooperative agents.

The diffusion protocol periodically refreshes reputation values maintained by each peer; this way, an extremely updated vision of other agent reputation is always maintained. Information obtained through the diffusion protocol contributes to determine the actual reputation value. The actual generic (good or bad) reputation value, held by peer i about peer j, $r_i^j(t)$, is computed adjusting the value collected so far, $r_i^j(t-1)$, with a weighted average of information obtained from other agents, $r_k^j(t-1)$. Weights used in this average are good reputation values of "gossipy" agents, $gr_i^k(t-1)$. The actual reputation value is computed according to equation (3).

$$r_i^j(t) = (1-\beta)*r_i^j(t-1)+\beta*\frac{\sum_{k \in K} gr_i^k(t-1)*r_k^j(t-1)}{\sum_{k \in K} gr_i^k(t-1)} \quad (3)$$

where K is the set of reliable gossipy agents for which the good reputation value exceeds an opportune threshold τ:

$$K = \{k : gr_i^k(t-1) \geq \tau\}. \quad (4)$$

A change of the threshold τ will affect peer willingness to rely on neighbours.

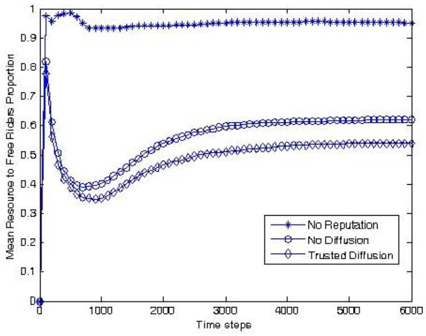

Fig. 1. Comparison between resources obtained by free riders and by cooperative peers

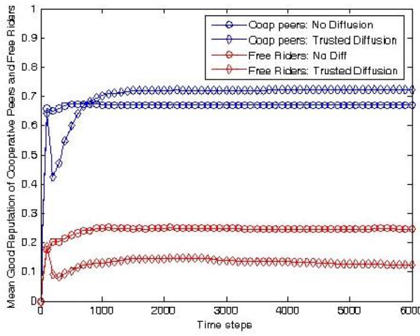

Fig. 2. Comparison between good reputation for free riders and for cooperative peers

V. Performance Analysis

In order to evaluate the proposed management reputation system we performed extensive testing in several simulation scenarios. For this purpose, a simple simulator has been developed and implemented in the java programming language, allowing modelling unstructured P2P networks. The aim of the simulator is to verify the effectiveness of the proposed reputation management system. In our simulations we randomly generated several network topologies, constituted by 100 peers. In our synthetic topologies each peer has a fixed number of neighbours. Neighbours are selected among all network peers according to a uniform distribution. For each network we assume that the 20% of peers are free riders, which provide only the 10% of requested services. A flooding-based algorithm was adopted to forward queries; queries concerned resources selected from a given domain, according to a uniform distribution. For each peer, the generation time between two successive queries follows a Pareto distribution. Using the same basic network scenario, we carried out simulations in three different configurations. In the first one, a very basic setting, any incentive mechanism is used, in the second one we introduced only the local reputation mechanism, while in the last one we adopted the full approach, constituted by the local reputation together with information diffusion.

The metrics used for comparisons were: the number of resources received by a free rider normalized with respect to the average number of resources received by a cooperative peer (Fig.1), and the mean value of actual good reputation for cooperative peers in comparison to good reputation of free riders (Fig.2). As shown in the Fig.1, without reputation management, a free rider has about the same resources than a cooperative peer. This is an expected result. The interesting result is that, using the local reputation, a free rider receives, on average, the 53% of resources with respect to those obtained by a cooperative peer. This percentage decrease up till the 44% when the reputation diffusion is enabled. In Fig.2 it is shown that the difference between good reputation for free riders and for cooperative peers is clear. Thanks to the reputation diffusion, this difference increases, allowing a more effective characterization of malicious peers.

VI. Conclusions and Future Works

Several works have been proposed in literature about reputation management system in distributed environments, but hardly any result is suitable for unstructured systems, in which there is no guarantee on cooperation from participating entities. In this paper, a totally distributed approach is proposed. Agents interact in order to autonomously determine the reputation value for agents with which they have interacted previously. Moreover, thanks to a reputation diffusion mechanism, attained by cluster of mutually reliable agents, the system is robust with respect to dynamic events, typically occurring in unstructured distributed environments. Our work also places the grounds for a system formalization through evolutionary game theory tools. This future development will allow to formally prove those properties that arose from experimental results.

References

[1] J. Hofbauer and K. Sigmund, "Evolutionary game dynamics," *Bull. Amer. Math. Soc.*, vol. 40, pp. 479–519, July 2003.

[2] M. Ripenau and I. Foster, "Mapping the gnutella network: Macroscopic properties of large-scale peer-to-peer systems," in *First International Workshop on Peer-to-Peer Systems*, 2002.

[3] S. Saroiu and P. Gummadi, "A measurement study of peer-to-peer file sharing systems," in *Multimedia Computing and Networking*, 2002.

[4] Epinions. [Online]. Available: http://www.epinions.com

[5] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks," in *12th International Conference on World Wide Web*, NY, USA, 2003, pp. 640–651.

[6] S. Songand, K. Hwang, R. Zhou, and Y. Kwok, "Trusted p2p transactions with fuzzy reputation aggregation," *IEEE Internet Computing*, vol. 9, pp. 24–34, nov-dec 2005.

[7] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins, "Propagation of trust and distrust," in *13th International World Wide Web Conference*, NY, USA, 2004, pp. 403–412.

[8] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster, "To share or not to share: An analysis of incentives to contribute in collaborative file sharing environments," in *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[9] B. Mortazavi and G. Kesidis, "Model and simulation study of a peer-to-peer game with a peputation-based incentive mechanism," in *Information Theory and Applications Workshop (IEEE)*, UC San Diego, USA, 2006.

[10] S. Kontogiannis and P. Spirakis, "The contribution of game theory to complex systems," in *10th Panhellenic Conference of Informatics (PCI)*, 2005, pp. 101– 111.

[11] C. Li, B. Yu, and K. Sycara, "An incentive mechanism for message relaying in peer-to-peer discovery," in *Second Workshop on the Economics of Peer-to-Peer Systems*, 2004.