This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE P2P 2011 proceedings.

SocialHelpers: Introducing Social Trust to Ameliorate Churn in P2P Reputation Systems

Marc Sánchez-Artigas and Blas Herrera
Department of Computer Engineering and Maths, Universitat Rovira i Virgili, Spain
Email:{marc.sanchez, blas.herrera}@urv.cat

Abstract— **Reputation systems rely on historical information to account for uncertainty about the intention of users to cooperate. In peer-to-peer (P2P) systems, however, accumulating experience tends to be slow due to the high rates of churn — the continuous process of arrival and departure of peers. The flow of transactions is continuously interrupted by departures, which can significantly affect the convergence of reputation systems. To shed light on this, this paper presents an accurate model for capturing the influence of churn on the process of building reputations. Using our model, system architects can determine the minimal transaction rate that guarantees fast convergence and design their systems accordingly. Unfortunately, the natural transaction rate of users is sometimes too low (e.g., due to physical constraints like network bandwidth, etc.) that many of them are likely to experience significant delays in the process of building reputations for their neighbors. We face this problem by leveraging the inherent trust in social networks. The basic idea is that users ask their social links to transact with strangers and together generate reputation ratings in a short time scale. Our simulation results report reductions of 50% or greater in the convergence time in environments with high churn rates.**

I. Introduction

Peer-to-peer (P2P) systems have become immensely popular in the Internet. Unfortunately, the open and anonymous nature of these systems often leads to a serious lack of accountability, opening the door to abuses by malicious peers. To fight against the misbehavior of peers, a plethora of reputation systems have appeared in the last few years with promising results [1], [2], [3], [4], [5], [6], to name a few. However, the existing literature mostly focuses on the vulnerabilities and countermeasures, and does not consider the significance of peer churn in the process of reputation building (see survey [7] and references therein). Any peer may be found interacting with a stranger at any time, a stranger that apparently offers a high quality service. As the stranger will probably leave and rejoin the overlay as any other user, the peer may be found waiting for the stranger to reenter the overlay again to accumulate enough evidence for trustable prediction.

Commonly, the performance of a reputation system has been measured by how accurately the calculated reputations predict the quality of future interactions. An optimal accuracy requires long-lived relationships and many transactions, which are hard to accomplish in an environment where users alternate between activity periods and OFF periods continuously. In P2P systems, one cannot expect that participating peers remain in the system for many consecutive transactions. As a result, it is very likely that any peer waits for several reconnections (ON/OFF cycles) of their neighbors before obtaining enough evidence to predict future behavior. Depending on the relative duration of the ON to OFF periods, the convergence time may be excessively long.

Another aspect that strongly determines the performance of a reputation system is the *natural transaction-rate limit*, or the maximum rate at which peers can interact. The transaction rate is typically determined by physical limitations such as network bandwidth and CPU, but very often depends on the application itself. For instance, many P2P applications impose upload rate limits to minimize uplink saturation [8]. Other applications use transaction rate limiters to limit the damage done by malicious users [9]. And even ISPs are known to rate-limit the bandwidth consumed by certain applications to reduce costs (the so-called Comcast/BitTorrent case is a nice example). All these methods lead to the reduction of the transaction rate and the subsequent increase in the convergence time. For many reputation systems, a low transaction rate can mean acting so slow that traitors can easily abuse the system. A traitor is a user that behaves well for some time to boost its reputation and later, quickly defects in many transactions.

To objectively determine the impact of churn on reputation building, we start this paper by developing a stochastic process to inspect whether, given a specific transaction rate, a moderate amount of churn limits the effectiveness of a reputation system. More concretely, our stochastic process describes the evolution of the number of interactions occurring between a consumer of a service and the provider of that service. To make this a worst case analysis, we assume that there is no historical information about that provider in the system. In such a case, the consumer must rely on direct observation as there is no indirect feedback available, which maximally interferes reputation building [10]. Under this analytical model, we then derive the expected delay until the consumer accumulates sufficient evidence to support either belief or disbelief towards that provider. Also, we obtain exact formulas for specific lifetime distributions and show that a high turnover rate can significantly limit the speed at which reputation is built over time.

However, the power of our analysis is that system architects can determine exactly the transaction rate that guarantees fast convergence and adapt their designs accordingly. This finding is very relevant for reputation systems where every entity runs a local instance of the reputation system to achieve scalability (see, for instance, [2], [11], [12] and [7] for more references). Indeed, we use our analytical tool to invent a social solution to

This research has been partially funded by the Spanish Ministry of Science and Innovation through project DELFIN, TIN2010-20140-C03-03.