Exploring the Feasibility of Reputation Systems under Churn

Marc Sánchez-Artigas, Pedro García-López, *Member, IEEE,* and Blas Herrera

***Abstract*—This letter examines how high rates of churn — the continuous process of node arrival and departure — affect rating mechanisms for peer-to-peer (P2P) networks. In particular, short peer lifetimes mean reputations are often generated from a small number of transactions, and thus are few reliable. To understand this relationship, this letter introduces an analytical model which determines the optimal transaction rate and the expected time to produce a reliable reputation, under both exponential and Pareto lifetime distributions.**

***Index Terms*—Reputation mechanisms, P2P networks, churn.**

I. Introduction

In recent years, a number of reputation systems have been proposed to protect peer-to-peer (P2P) networks against the increasing misbehavior of peers. For instance, in a file-sharing network, selfish users may decide to act in their best interests not offering their resources for sharing. Reputations that reflect the perceptions of individual peers can deter moral hazard and discourage abuse.

Unfortunately, existing reputation mechanisms do not cope well with the dynamic nature of P2P networks. As reputations are built using historical feedback over a relatively long period of time, shorter peer lifetimes imply very often that ratings are inaccurate. For applications that rely on peers for data storage, message forwarding or distributed computing, choosing a peer based on short-term reputations is highly undesirable. One way to address this shortcoming is to use proactive reputations [6].

Proactive reputations augment traditional reputation systems with on-demand transaction feedback. The key idea is to allow peers to proactively initiate transactions for the sole intention of generating reliable reputations for new peers or those peers with short lifetimes.

However, proactive reputations require that peers determine under which rates of churn ratings should be proactively built. Otherwise, the additional traffic injected into the system might drain away the benefits brought by reputation.

To examine this issue, in this letter we propose a stochastic model to understand the behavior of reputation systems under churn. For this aim, we define two new metrics: the probability of reliable reputation generation and the mean generation time. The former represents the sensitivity of a reputation system to churn, while the latter measures its responsiveness. With these metrics, we clarify why the distribution of the residual lifetime plays an important role in the feasibility of reputation systems.

In general, there exists a lack of systematic studies, specially theoretical analysis, on evaluating how well reputation systems tolerate churn. With this letter, we hope to start filling this gap.

II. Peer-to-Peer Network Model

In this letter, we consider that each joining peer has several neighbors whose behavior is analyzed by means of a reputation system. These neighbors provide some services (such as data storage, message forwarding or distributed computing) that the joining peer consumes. Thereby, the purpose of the reputation system is to predict the future behavior of each neighbor based on past transactions, for it needs reliable ratings.

Similar to [4], we consider that each joining peer is assigned a random lifetime drawn from some distribution $F(x)$, which reflects the amount of time the peer stays in the network. Since it is almost impossible for a joining peer to pick its neighbors such that their arrival times are exactly the same as itself, it is plausible to assume that neighbor selection is performed over the peers already present in the system. These peers have each been alive for some random amount of time, which means that the residuals of their lifetimes has to be considered rather than their lifetimes.

Denote by R the remaining lifetime of a neighbor when the joining peer entered the system. Assuming that the system has reached stationarity, the cumulative density function (CDF) of residual lifetime is given by [5]:

$$F_R(x) = \Pr(R < x) = \frac{1}{E[L]} \int_0^x (1 - F(z))dz \qquad (1)$$

where $F(x)$ is the CDF of peer lifetime L, and $E[L]$ its mean. As in [4], we will experiment with two lifetime distributions:

- *Exponential lifetime*, with CDF $F(x) = 1 - e^{-\lambda x}$, $\lambda > 0$; and
- *Pareto lifetime*, which has been reported by several works to provide a tight fit to the real lifetime distribution found in P2P systems. Its CDF is $F(x) = 1 - (1 + \frac{x}{\beta})^{-\alpha}$, $\alpha > 1$, where parameter α represents the heavy-tailed degree and β is a scale parameter.

For exponential lifetimes, the residual lifetimes are trivially exponential thanks to the memoryless property of $F(x)$, i.e., $F_R(x) = 1 - e^{-\lambda x}$. However, for Pareto lifetimes, the residuals are *more heavy-tailed* and exhibit shape parameter $\alpha - 1$, i.e., $F_R(x) = 1 - (1 + \frac{x}{\beta})^{1-\alpha}$, $\alpha > 1$.

It has been shown that average request rates of peers do not change significantly over time in BitTorrent [3], enabling the mathematical treatment of query arrivals as a Poisson process. Here, we also consider that the number of transactions made

Manuscript received March 10, 2009. The associate editor coordinating the review of this letter and approving it for publication was G. Lazarou.

The authors are with the Department of Computer Engineering and Mathematics, Universitat Rovira i Virgili, Tarragona, 43007 Spain (e-mail: marc.sanchez@urv.cat).

This work is supported through research project P2PGRID, no. TIN2007-68050-C03-03, of the Ministry of Education and Science, Spain.

Digital Object Identifier 10.1109/LCOMM.2009.090573