Activity Stereotypes, or How to Cope with Disconnection during Trust Bootstrapping

Marc Sánchez-Artigas, *Member, IEEE,* and Blas Herrera

Abstract—Trust-based systems have been proposed as means to fight against malicious agents in peer-to-peer networks, volunteer and grid computing systems, among others. However, there still exist some issues that have been generally overlooked in the literature. One of them is the question of whether punishing disconnecting agents is effective. In this paper, we investigate this question for these initial cases where prior direct and reputational evidence is unavailable, what is referred in the literature as *trust bootstrapping*. First, we demonstrate that there is not a universally optimal penalty for disconnection and that the effectiveness of this punishment is markedly dependent on the uptime and downtime session lengths. Second, to minimize the effects of an improper selection of the disconnection penalty, we propose to incorporate predictions into the trust bootstrapping process. These predictions based on the current activity of the agents shorten the trust bootstrapping time when direct and reputational information is lacking.

1 Introduction

Trust-based systems have been proposed for a large variety of applications, ranging from mobile ad-hoc networks, Grids and P2P networks. At present, despite their maturity, some fundamental questions have still left unanswered. One of these important questions relates to the notion of disconnection as a *trust diminishing* event or *punishable* action. In open environments like P2P systems and Grid platforms like BOINC [1], disconnection affects the quality of service (QoS). To wit, in a P2P streaming service, QoS can be achieved as long as a continuous and uninterrupted data flow is maintained. It is basically for this reason that streaming systems like ripple-stream [2] and trust systems like [3], [4], [5] issue negative feedback for agents that are supposed to be providing the service but cannot do so because they are disconnected. The key problem is that in open environments it is not possible to differentiate between negative feedback due to malicious behavior and negative feedback due to disconnection; an agent can disconnect at any time and the trustor cannot tell whether the disconnection was intentional or not.

By the above discussion, one could infer that the most convenient method is to heavily penalize disconnection. However, contrary to intuition, as we show in this work, punishing disconnection might be *counterproductive*. This is especially true in those initial situations where no prior direct and reputational evidence is available. One case is when a new agent enters the system for the first time. In this situation, it is generally not possible for any trustor to form a reliable opinion on that agent. This also occurs when users form an ad-hoc group around a shared goal and disband once the pursued goal is met. In such cases, evidence can only be obtained through direct interaction, when some trustors are willing to take a chance and risk interacting with unknown trustees. It is the risk inherent in bootstrapping trust that can lead to applying little or no penalty on disconnecting agents.

For instance, consider the case that multiple unknown agents offer the same file to download. In this context, a transaction may simply be the transfer of a file piece to a trustor. Since a priori all agents have the same unknown disposition to good action, the first interacting trustee is chosen at random. Now suppose that after completing a certain number of transactions, the interacting trustee is unresponsive. At that point, the trustee may accumulate a certain amount of positive feedback and present a good trust level. Then it is not hard to imagine that the trustor gets confronted with the decision of whether to wait for the trustee to recover or take a chance on another agent.

The magnitude of the penalty determines the outcome of that decision. If the penalty is large, the odds to take a chance on a new agent are higher. In this case the trustor will maximize interaction. But it will be more exposed to abuses. On the contrary, if the penalty is low, the trustee may come online before getting low trustworthiness and continue providing good service. This will minimize the risk of bad interaction but at the expense of more service interruptions.

Our first contribution is to examine this trade-off, and more generally, to assess to which extent the amount of penalty given to disconnection affects the bootstrapping of trust. To make the analysis tractable, we assume that T time units must elapse after disconnection in order to prefer an unknown agent. A smaller value of T implies a greater penalty, i.e., a higher probability for the trustor to take a chance on a new trustee. Using this parameter, we develop a stochastic model to estimate the expected time to obtain the first confident trust evaluation on an agent, provided that all the agents implementing the service are unknown to the trustor. By "confidence" we refer to the event that the trustor acquires enough direct evidence to

- *Department of Computer Engineering and Maths, Universitat Rovira i Virgili, Spain. Email:*`{marc.sanchez, blas.herrera}@urv.cat`

form a risk perception. This time, we simply called it the *trust bootstrapping time*, is a good indicator of the efficacy of a trust system. If this time is short, trustors can quickly form an impression to guide their future interactions.

As a result of our analysis, we arrive at the conclusion that *there is not a universally optimal penalty*. The optimal penalty is too much dependent on the exact amount and type of churn, or the term coined to refer to the collective effect of the continuous arrival and departure of users in the system [6]. The direct consequence of this is that an unfortunate choice for the penalty given to disconnection can lengthen the trust bootstrapping time. To the best of our knowledge, we are the first to analyze the effect that disconnection punishment has on trust bootstrapping.

Our second contribution is to incorporate availability predictions into the trust bootstrapping process to reduce the effects of a bad selection of the disconnection penalty. The key feature of our predictions is that they are based on the activities of users. Returning back to our example, a trustor may simply learn that the trustees downloading files between 1-4 GBs tend to have long sessions, and use this knowledge to choose between the candidates based on their downloading activity. We call these predictions *activity stereotypes*. While one can argue that this concept is similar to the "classical" concept of stereotypes [7], [8], the specifics of user connection habits require specialized treatment, making activity stereotypes a novel approach.

The remaining of the paper is structured as follows. In Section 2, we overview related work. We introduce our analytical model of trust bootstrapping and dynamics in Section 3. Section 4 describes the analytical results that demonstrate the lack of a universally optimal penalty for disconnected operation. Section 5 describes the notion of activity stereotypes and Section 6 provides an evaluation of their effectiveness. Conclusions are drawn in Section 7.

2 Related Work

In general, the differentiation between a malicious and a good user depends heavily on the nature of the system. Misbehavior can be classified as deliberately malicious or arise out of temporary outages or user connection habits.

Deciding whether to qualify such *non-subjective* factors as misbehavior depends on the parameters of the system. While many works on trust, either implicitly or explicitly, have formerly classified *'No Response'* and disconnection as a punishable event [2], [3], [4], [9], to name a few, the consequences of punishing disconnected behavior have not received sufficient attention. Typically, trust systems like PET [3] and ripple-stream [2] fix a numeric constant to penalize disconnection, thereby punishing agents in a way completely irrespective of their connection patterns. As we demonstrate in this work, this may be problematic when interacting with strangers, and in general in those cases where both direct and reputational evidence is not forthcoming. The present article is the first to investigate this important issue by studying the effects of punishing disconnecting agents and proving the lack of an optimal penalty for disconnection.

The general effects of dynamics on trust systems have been studied in our prior works [10], [11], [12]. In these works, we developed an analytical framework to assess the difficulties of establishing trust in the absence of trust information sharing among entities. Thus, no analysis of disconnection punishment was conducted, and even less the development of a new technique to bootstrap trust.

For trust bootstrapping, recently there have been some efforts to develop tentative forms of trust in the absence of direct and reputational experiences [7], [8], [13]. These approaches propose to exploit stereotypical impressions formed on "similar" agents in previous contexts in order to make tentative trust evaluations on unknown agents. Although the concept of stereotypes for decision making was proposed previously (see, for instance, [14]), the first computational stereotype model was introduced by Liu et al. in [7]. Since then, other works like [8], [13] proposed to exploit stereotypes for trust evaluation. In particular, Burnett et al. [8] proposed to bootstrap trust of unknown agents through stereotypes, which are built based on the M5 tree learning algorithm and shareable as reputations.

The principal idea of the above described stereotypical approaches is to utilize visible features of agents to make generalized trust assessments. In this regard, "classical" stereotypes resemble activity stereotypes. However, they present two important differences. The first difference is that activity stereotypes focus on *service continuity* rather than on the *trustworthiness* of agents, as occurs in the case of classical stereotypes. The other basic difference is that activity stereotypes are always formed from information available in the system, while classical stereotypes can be built with featural evidence coming from *external* sources of information like social networking systems. To inform this argument, Burnett et al. [15] have recently discussed what types of contextual knowledge can be used to build stereotypes. Three feature sources are identified, but any of them can be directly used to build activity stereotypes, since they correlate with trustworthiness instead of with service continuity, like the relationships between agents in social networks and the experience of agents in certain tasks. As a result, activity stereotypes require a different and original treatment compared with prior approaches.

Finally, Sensoy et al. [16] argue that agents exhibiting similar behavior share some patterns in the relationships between their descriptive features and propose to exploit them to bootstrap trust. As the above proposals, the main flaw of this approach is the use of ontological knowledge to describe agents in detail, information which is seldom available in many scenarios. Like us, they use Subjective Logic to represent trust [17] and the base rate to integrate their predictions into the trust formation process as [8].

3 Modeling Trust Bootstrapping

3.1 Agent Dynamics and Metrics

As in many important modeling works [18], [19], [20], we model the alternating online and offline agent behavior as a 2-state continuous-time Markov chain (CTMC) with transition rates r_{on} and r_{off}, respectively. That is, a user stays connected for an average amount of time $1/r_{on}$ and then disconnects. The average amount of time the agent stays offline before reconnecting to the system is $1/r_{off}$. In systems dominated by user-driven interruptions like Maze or Kad, it has been shown that this simple CTMC provides a good approximation to user behavior [21].

Having described the model for agent churn, we now turn to *trust bootstrapping*. Since we consider initial cases where direct and reputational information is unavailable, initial trust can only be obtained from direct encounters. It may be unwise to aggregate the opinions of unknown agents, as some may be malicious or unreliable in some way. This is the fundamental reason why we characterize trust bootstrapping as a stochastic process modeling the occurrence of direct interactions. Concretely, we assume that a trustor must interact at least ℓ times with a trustee to be *confident* that the resulting trust evaluation is a good predictor of future behavior. That is, the value of ℓ marks the point at which there is little or no uncertainty about the outcome of the next interaction, and when it is safe to ask reputational queries to a trustee, among other things.

Further, the use of this parameter turns our study into a generic analysis. That is, while the exact value of ℓ will vary across systems, our approach will remain valid for many of them. For instance, this includes those trust and reputation systems that combine the numerical ratings of past interactions to output a trust value, including those based on belief models [17] and all the cited works in the related work like [2], [3], [4], [9], [8], [16], to name a few.

Based on this parameter, we define a new metric called the *trust bootstrapping time* to understand the influence of punishing disconnection when interacting with *unknown* agents. This measure focuses on the time it can take for a trustor within a group of unknown agents to form the first *confident* trust evaluation. Ideally, this time should be as short as possible so that the trustor could benefit from confident trust evaluations in their partners before group disbandment.

Definition 1. *Given a group of unknown agents $\mathcal{U}$ providing a service, we define the trust bootstrapping time τ_ℓ as the time to complete the first ℓ transactions with any of the agents in $\mathcal{U}$.*

For analytical tractability, we assume that transactions occur *immediately one after another*, according to a Poisson process with parameter λ. In this way, transactions have a similar duration, thus making unnecessary to calculate the gain in trust on the basis to the amount of work done. Also, this interaction model is very common in the trust literature (see [22], [10], to cite a few examples).

3.2 Stochastic Model

Since all agents in the group are new and unknown, all are assigned the same *default* trust value. In practice, this value represents how much trust the trustor places in an agent before any evidence has been received. Because all agents have the same trust value, we assume that trustor chooses one agent uniformly at random to interact with. At that moment, the trust bootstrapping process starts.

For analytical tractability, we assume that the outcome of a transaction is *always* satisfactory to initiate a new one with the current trustee. Therefore, our stochastic model considers only one type of trust decreasing event: the '*No response*', which occurs when the trustee, intentionally or not, fails to complete a transaction due to disconnection. This corresponds to the case where agents provide good service but receive punishment due to their online-offline oscillatory behavior, the subject of study of this work. We observe that this is the right way to isolate the effects of dynamics from the effects of malicious behavior in trust bootstrapping, which has been already studied [7], [8].

Also, it must be noted that the effect of this assumption is not significant. We notice that if trustees behave badly by responding *wrongly* or even *maliciously* to transaction requests, the trust bootstrapping time will be longer. The reason is that *it generally takes a few bad interactions to lose trust but many good interactions to trust someone* [23]. Then, it is natural to expect that the trustor switches to a new agent after experiencing a just few negative transactions, lengthening the time to complete the first ℓ transactions with some agent in the group. In practice, this makes our results conservative but accurate enough to measure the effect of disconnection. Indeed, our analytical predictions are in good agreement with our simulation results where half of the trustees were instrumented to misbehave. See Section 6 for further details.

Last but not least, there remains the question of how to punish disconnecting agents *stochastically* speaking. To characterize the intensity of the penalty, we introduce a new parameter into the model. This parameter, denoted by T, specifies the number of time units needed to elapse after disconnection of the current trustee to take a chance on an unknown agent in the group:

Definition 2. *A switch to an unknown agent is triggered after the disconnection of the current trustee during T, $T > 0$, time units. During this time, the disconnected trustee is assumed to reject all attempted transactions issued by the trustor.*

By this characterization, it is possible to make sure that an unknown agent is always preferred over the trustees who had been disconnected for more than T consecutive time units, in an attempt to maximize service continuity.

Further, the value of T determines the *aggressiveness* of the punishment. Smaller values of T represent a greater penalty, i.e., a higher probability for the trustor to take a chance on a new agent. On the contrary, larger values of T decrease the risk of bad response. Considering that the current trustee is offering good service, a larger T trades off longer interruptions in the service against the risk of switching to an unknown agent, who may be malicious. By simply varying the value of T, our model allows us to investigate how the intensity of disconnecting penalties impacts the trust bootstrapping time.

The state transition diagram is given in Fig. 1. States (ON, i) and (OFF, i) represent the case where the number of completed transactions is $i \geq 0$ and the current trustee is ON and OFF, respectively. In state (ON, i), the process

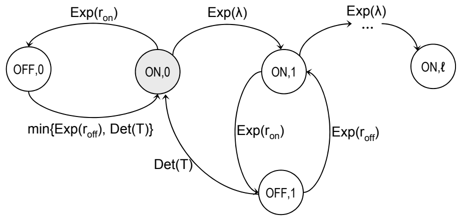

Fig. 1. State diagram for semi-Markov chain $\{Y(t)\}_{t\geq 0}$.

can jump into either state (ON, $i+1$), which represents that a new transaction has ended, or state (OFF, i), which implies that the trustee is now offline. In this state, the process can jump into state (ON, 0) if disconnection time exceeds T, which implies a *trustee switch*, or to state (ON, i) otherwise. The state of the process at time 0 is of course (ON, 0). The kernel $\mathbf{Q}(t) = \left[Q^{\xi,i}_{\upsilon,j}(t)\right]$ of this process we denote by $\{Y(t)\}_{t\geq 0}$ is as follows:

$$\begin{aligned}
Q^{\text{OFF},0}_{\text{ON},0}(t) &= \Pr\{\textit{trustee recovers before or at time } t,\ t<T,\\
&\quad \textit{with no transactions completed}\}\\
&= 1-e^{-r_{off}t}+e^{-r_{off}t}u(t-T).\\
Q^{\text{ON},i}_{\text{OFF},i}(t) &= \Pr\{\textit{trustee logs off during transaction } i+1\}\\
&= \tfrac{r_{on}}{r_{on}+\lambda}\left(1-e^{-(r_{on}+\lambda)t}\right),\quad \forall i=0,...,\ell-1.\\
Q^{\text{ON},i}_{\text{ON},i+1}(t) &= \Pr\{\textit{transaction } i+1 \textit{ completed before the}\\
&\quad \textit{trustee goes to OFF state}\}\\
&= \tfrac{\lambda}{r_{on}+\lambda}\left(1-e^{-(r_{on}+\lambda)t}\right),\quad \forall i=0,...,\ell-1.\\
Q^{\text{OFF},i}_{\text{ON},i}(t) &= \Pr\{\textit{trustee goes to ON state before}\\
&\quad \textit{or at time } t \textit{ and } t<T\}\\
&= 1-e^{-r_{off}t}-\left(e^{-r_{off}L}-e^{-r_{off}t}\right)u(t-T),\\
&\quad \forall i=1,...,\ell-1.\\
Q^{\text{OFF},i}_{\text{ON},0}(t) &= \Pr\{\textit{trustee does not go to ON state}\\
&\quad \textit{at time } t \textit{ and } t\geq T\}\\
&= e^{-r_{off}T}u(t-T),\quad \forall i=1,...,\ell-1.
\end{aligned}$$

where $u(t-T)$ is the unit step function at T.

Because we are interested in the time to complete the first ℓ transactions with any of the unknown trustees, it can be easily verified that the trust bootstrapping time τ_ℓ corresponds to the first-hitting time of process $\{Y(t)\}_{t\geq 0}$ onto state (ON, ℓ) given that $Y(0)=(\text{ON},0)$:

$$\tau_\ell = \inf\{u>0 : Y(u)=(\text{ON},\ell)|Y(0)=(\text{ON},0)\}.$$

In the next section, we calculate the expectation of the first-hitting time and use it to observe what happens to the trust bootstrapping time as function of the intensity of the penalty imposed on disconnection.

4 Disconnection Punishment: Analysis

We start by finding the average bootstrapping time $\mathbb{E}[\tau_\ell]$ and measure the influence of disconnection punishment on trust evaluation when no prior evidence can be found:

Theorem 1. For user ontimes with CDF $1-e^{-r_{on}x}$, user offtimes with CDF $1-e^{-r_{off}x}$ and punishment intensity T, the mean time to complete the first ℓ transactions with a trustee in the group is given by:

$$\begin{aligned}
\mathbb{E}[\tau_\ell] &= \frac{1}{r_{off}\lambda^\ell}\sum_{i=0}^{\ell}\left(r_{on}e^{-Tr_{off}}\right)^{\ell-i}S_{i,\ell}, \qquad (1)\\
S_{i,\ell} &= \left(\eta_{\ell,i}r_{off}\lambda^{i-1}-\epsilon_{\ell,i}\lambda^i+\epsilon_{\ell,i}\lambda^i e^{Tr_{off}}\right),
\end{aligned}$$

where $r_{on}=1/\mathbb{E}[L]$, $r_{off}=1/\mathbb{E}[D]$, and $\mathbb{E}[L]$ and $\mathbb{E}[D]$ denote the average online and offline session durations, respectively. Further, $\eta_{\ell,i}$ and $\epsilon_{\ell,i}$ satisfy the recurrence relations: $\eta_{\ell,i}=\eta_{\ell-1,i}+\eta_{\ell-1,i-1}$ for all $i>1$, $\epsilon_{\ell,i}=\eta_{\ell,i+1}$ for all $i<\ell$. The initial conditions are: $\eta_{\ell,0}=0$, $\eta_{\ell,1}=1$, $\eta_{\ell,\ell}=\ell$ and $\epsilon_{\ell,\ell}=0$.

Proof: The proof has been deferred to Appendix A. □

The first observation to be made is that the pure effect of the alternating online-offline behavior of agents can be predicted by taking Eq. (1) to the limit ($T\longrightarrow\infty$). Recall that letting $T\rightarrow\infty$ is equivalent to imposing no penalty on disconnected agents. Hence, the trustor considers that disconnection is temporary, e.g., due to the breakdown of the Internet connection, and it is worthwhile waiting for the trustee to recover. By taking $T\longrightarrow\infty$, we obtain the following corollary from Theorem 1:

Corollary 1. When imposing no penalty on disconnected agents, the mean bootstrapping time for ℓ transactions is given by:

$$\mathbb{E}[\tau_\ell]=\frac{\ell}{\lambda}\left(\frac{r_{off}+r_{on}}{r_{off}}\right), \qquad (2)$$

where λ is the transaction rate, and r_{on} and r_{off} denote the disconnection and reconnection rates, respectively.

Proof: A rigorous proof is given in Appendix B. □

Eq. (2) has a very interesting interpretation: the mean bootstrapping time is *inversely proportional* to the steady-state user availability $\mathcal{A}=\frac{\mathbb{E}[L]}{\mathbb{E}[L]+\mathbb{E}[D]}=\frac{r_{off}}{r_{on}+r_{off}}$.

Because there is no punishment in this case ($T\rightarrow\infty$), Eq. (2) suggests that the time spent in accruing enough supporting evidence to make a confident trust evaluation depends basically on the probability of the initial trustee to stay connected. This carries consequences for trustors who seek to minimize the inherent risk in bootstrapping trust. A cautious trustor will prefer to impose little or no penalty to minimize the number of switches to unknown agents caused by the temporary outages of a cooperative trustee. Such a "conservative" behavior will cause trust evaluations *to be highly dependable on agent availabilities*. If availabilities are low, it will be clearly advantageous to select another agent after a short period of inactivity with no way to reduce the inherent risk in bootstrapping trust. This result is pessimistic for certain types of distributed systems like P2P networks. In P2P systems, availabilities tend to be low. Hence, imposing no penalty may convert trust bootstrapping into a rather lengthy process. Taking, for instance, the mean availability observed in BitTorrent

[24], the average bootstrapping time $\mathbb{E}\left[\tau_\ell\right]$ will be of 3.57 hours for $\ell = 10$ and transaction rate λ of 10 interactions per hour. More examples are reported in Appendix C.

4.1 Impact of Disconnection Penalty

As just elaborated above, it seems at first glance that an aggressive punishment should shorten the bootstrapping time and favor a more uninterrupted service. While the latter is true since the periods of inactivity are shorter, the former is not necessarily so. Contrary to intuition, a short T may increase the trust bootstrapping time, specially if the extended design principle that good behavior should increase trust slowly is applied. In PET [3], for example, the penalty incurred by not responding is 3 times greater than the reward obtained for good behavior.

An example of this behavior is illustrated in Fig. 2(a). In the figure, the predicted $\mathbb{E}\left[\tau_\ell\right]$ is plotted against ℓ and compared to the empirical $\mathbb{E}\left[\tau_\ell\right]$ obtained via simulation. For this example, the average ontime was $\mathbb{E}\left[L\right] = 1$ hour, the average offtime $\mathbb{E}\left[D\right] = 0.5$ hours with a transaction rate λ of 10 transactions per hour. As seen in the figure, as the number of interactions ℓ needed to form a reliable opinion increases, the average bootstrapping time under an aggressive punishment ($T = 6$ minutes [1]) may end up becoming greater than when no punishment is imposed on departing agents ($T = \infty$). The primary reason is that disconnected agents present a strong tendency to return sooner than later as observed in real systems [6]. Hence, an aggressive strategy to support service continuity may be counterproductive for trust establishment. In fact, the probability of switching to a new and unknown agent at least once is given by:

$$\left(1 - \rho^\ell\right) e^{-r_{off} T}, \quad \text{with } \rho = \frac{\lambda}{\lambda + r_{on}}, \tag{3}$$

which is dominated by ρ^ℓ as $T \to 0$ (as the punishment intensity level is raised). To put in perspective, consider that the number of transactions required to form a useful trust opinion is $\ell = 10$. In this scenario, Eq. (3) yields a probability of switching of 0.61 when $T = 0$ (immediate switch after disconnection), which is reduced to 0.41 for $T = 12$ minutes, and to 0 when $T \to \infty$ (no penalty).

Moreover, this numerical example clearly captures the existing trade-off between quality of service and the risk of interaction: optimizing service continuity may require, in more or less degree, to take a chance on an unknown, distrusted agent, which, in addition to increasing the risk associated with interaction, may not help to shorten the trust bootstrapping time. The same behavior can be seen for other choices of parameters as illustrated in Fig. 2(b). As can be seen in the figure, there exists a crossover point ($\ell = 14$ in this case) beyond which imposing no penalty on departed agents is again more effective.

In fact, if we concentrate on the amount of punishment T that minimizes the mean bootstrapping time, it can be shown that there exits an integer L with the property that

$$\lim_{T \to \infty} \mathbb{E}\left[\tau_\ell\right] < \mathbb{E}\left[\tau_\ell\right](0), \, \forall \ell \geq L \tag{4}$$

where $\mathbb{E}\left[\tau_\ell\right](0)$ denotes the average bootstrapping time obtained when $T = 0$ (maximum penalty), i.e., when the trustor switches immediately to an unknown agent after the disconnection of the actual trustee. Loosely speaking, (4) tell us that there exists a threshold L on the number of transactions beyond which the imposition of no penalty at all leads to a shorter bootstrapping time than the one reached when imposing an aggressive penalty, which is counterintuitive. Exact conditions of when this transition happens are given in Appendix D.

To wit, for the same choice of parameters of Fig. 2(a), the value of L is 9, which marks the point beyond which punishing disconnection becomes negative. This is better reflected in Fig. 2(c) where the mean bootstrapping time is plotted for different punishment intensities. The main point to be made here is the presence of two *differentiated* regions in the figure, separated by the crossover point L. For $\ell = 1, \ldots, L-1$, the minimum bootstrapping time is achieved when the penalty is maximal ($T = 0$), while for $\ell \geq L$, the optimal bootstrapping time is attained when no penalty is applied. Intermediate values strike balance between quick bootstrapping and service continuity.

In summary, our analysis proves *the lack of a universally "optimal" penalty for disconnection* and *formally* shows that *the effectiveness of the punishment is strongly dependent on the type and amount of agent turnover*.

5 Activity Stereotypes

The lack of a *global optimal disconnecting penalty* demands new techniques to improve bootstrapping of trust while protecting the system from high agent turnover.

One way to do so is to incorporate predictions on peer uptimes into the trust bootstrapping process. Among the possible solutions, we concentrate on the current activity of an agent as a *predictive* mechanism. By ascribing initial peer selection to learned classes of *agent activity*, a trustor can make use of prior experience in these classes to form tentative evaluations on the dependability component of trust, which is critical in the initial cases discussed in this article. This concept is similar to the notion of "classical" stereotypes [7], [8] but targeted at reducing the impact of disconnections and service interruptions in initial cases where prior evidence is unavailable.

To put it in a practical context, typical activities can be the download of a file in *pieces* (Direct Connect, Gnutella, ...); the participation in a live streaming session (PPLive, PPStream, ...). Concrete examples of activities and their respective *stereotypes* are discussed in Appendix E.

A requirement for activity stereotypes is that they are meant to complement, not replace, direct evidence about an agent when it is available. While activity stereotypes may facilitate useful predictions in initial cases, they are based on empirical generalizations, and should carry less weight than direct observation. Like in [8], we fulfill this

[1] Note that $T = 6$ minutes is equivalent to the average duration of a transaction: $1/\lambda$, for $\lambda = 10$ interactions per hour.

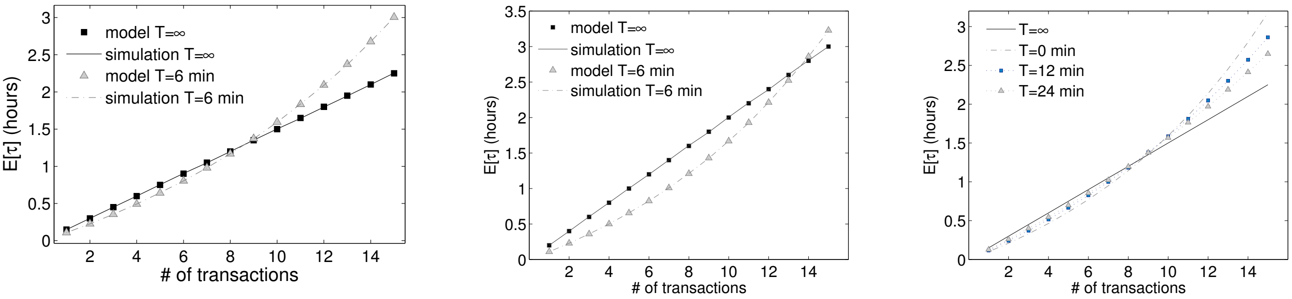

(a) $\mathbb{E}[L] = 1$ hour, $\mathbb{E}[D] = 0.5$ hours.

(b) $\mathbb{E}[L] = 1$ hour, $\mathbb{E}[D] = 1$ hour.

(c) $\mathbb{E}[L] = 1$ hour, $\mathbb{E}[D] = 0.5$ hours.

Fig. 2. $\mathbb{E}[\tau_\ell]$ plotted against ℓ for different mean ontime $\mathbb{E}[L]$ and offtime $\mathbb{E}[D]$ durations. $\lambda = 10$ transactions/hour.

TABLE 1
Trust model main notation.

Notation	Description
w_y^x	Opinion of a trustor x about an agent y
$P\left(w_y^x\right)$	Trust value for agent y derived from w_y^x
$\omega_{\mathbb{B}}$	Penalty for bad transactions
$\omega_{\mathbb{N}}$	Penalty for *'No response'*
ℓ	Number of transactions to bootstrap trust
λ	Transaction rate
$\vec{A}_y$	Activity vector of an agent y
$f(\vec{A}_y)$	Stereotypical trust value for agent y
a_{def}	Default trust value without stereotypes
a_{max}	Threshold on stereotypical trust values

requirement by adapting *default trust* in unknown agents to the behavior of the majority of agents carrying out the same activity. Recall that default trust refers to the trust put in a newly deployed trustee before any evidence has been received [25]. If the agents participating in a given activity A_i have historically tended to stay connected for long periods of time, an unknown agent executing A_i can be privileged by assigning a high initial trust to it.

Technically, the objective of our approach is to identify a function f that maps the activity vector of an agent $\vec{A}$ to an estimate $f(\vec{A})$ on service continuity which increases or decreases the default trust given to an unknown agent.

5.1 Trust Model

Regardless of the underlying model, the key requirement of our approach is that the estimates function f produces are compatible with the trust model being used. For this reason, we describe next a concrete trust model to show how easy it is to integrate activity stereotypes into a trust system. The primary requirement is that the trust system allows to assign neutral or default trust values to newly joined agents, which holds in the majority of cases.

Specifically, we adopt the model proposed in [17] and based on *Subjective Logic*. The reason is that it admits the mapping of the estimates from activity stereotypes onto default trust values, referred to as base rates in this trust model. Deterring participation by those individuals who are dishonest and encouraging trustworthy behavior are matters of the underlying trust system. We have simply augmented the trust model proposed in [17] with activity stereotypes to demonstrate their effectiveness. Any trust model using numerical ratings could be used in its stead.

5.1.1 Representation

In this trust model, an *opinion* held by a trustor x about a trustee y is represented as a tuple $w_y^x = \langle b_y^x, d_y^x, u_y^x, a_y^x \rangle$, where values b_y^x, d_y^x, and u_y^x express the degree of belief, disbelief, and uncertainty towards y, respectively. These values satisfy $b_y^x + d_y^x + u_y^x = 1$, with b_y^x, d_y^x, $u_y^x \in [0, 1]$. Uncertainty measures the absence of evidence to support either belief or disbelief, such that an opinion based on $\ell = 100$ transactions has a greater certainty than another one based on just 1 observation. Clearly, certainty or the confidence on a trust value is thus equivalent to $(1-u_y^x)$.

Worthy of special mention is the parameter $a_y^x \in [0; 1]$, called the *base rate*, which corresponds to *default trust*. In the absence of any specific evidence about a given agent, the base rate determines the a priori trust that would be put in any member of the group. For instance, if a_y^x is set to 0.75, we believe that the result of the first interaction with agent y will likely to be favorable.

5.1.2 Evidence Aggregation

Opinions are formed upon the basis of evidence amassed by interacting with other agents, which is represented as observed frequencies of positive and negative outcomes. A body of evidence held by a trustor x is a pair $\langle r_y^x, s_y^x \rangle$, where r_y^x is the number of positive transactions received from y, and s_y^x is the number of negative experiences. Using these parameters, an opinion is produced as [17]:

$$\begin{aligned} b_y^x &= r_y^x / \big(r_y^x + s_y^x + 2\big), \quad d_y^x = s_y^x / \big(r_y^x + s_y^x + 2\big), \\ u_y^x &= 2 / \big(r_y^x + s_y^x + 2\big). \end{aligned} \tag{5}$$

Notice that Eq. (5) guarantees that uncertainty decreases as more evidence is accumulated. Alternatively, evidence could be obtained from third parties who had interacted with this specific individual before. However, since we are examining the problem of establishing trust when no historical information is available, evidence is acquired first hand by each trustor.

To treat *no responses* as a bad action like in PET [3] and [4], so that the users who continuously join and leave the system receive low trustworthiness, we classify negative experiences in two different categories: '*No Response*' and '*Bad Behavior*'. '*No Response*' describes the situation where a trustee rejects or fails to complete a transaction request due to disconnection. Because it cannot be distinguished whether a disconnection was intentional or not, this class includes both cases. The reason is that, irrespective of the cause of the disconnection, the result is the same: a *service interruption*. The term '*Bad Behavior*' encapsulates all the other negative experiences, mostly malicious responses.

To materialize this distinction, we calculate s_y^x as linear combination of the observed frequencies of each type of outcome: $s_y^x = \omega_{\mathbb{B}} \cdot s_{y:\mathbb{B}}^x + \omega_{\mathbb{N}} \cdot s_{y:\mathbb{N}}^x$, where $s_{y:\mathbb{B}}^x$ is the number of wrong transactions, $s_{y:\mathbb{N}}^x$ is the number of transactions that received no response, and $\omega_{\mathbb{B}}$ and $\omega_{\mathbb{N}}$ are the weights attached to each type of negative action. These weights are used to assign different levels of importance to each type of negative experience.

5.1.3 Trust Metric

A single-valued trust metric, useful for ranking agents, can be derived from a specific opinion w_y^x as follows [17]:

$$P\left(w_y^x\right) = b_y^x + a_y^x \cdot u_y^x, \tag{6}$$

where the trust value is the probability expectation value $P\left(w_y^x\right)$ calculated from w_y^x. Observe that the base rate a_y^x determines the effect that the parameter u_y^x will have on the resultant trust value.

In many trust systems, the default value of the base rate is usually 0.5, which signals that before any evidence have been received, both outcomes are equally likely to occur. Also in this case, $P\left(w_y^x\right) = 0.5$, which is the least informative value about an agent when no evidence have been acquired. Further, in this case uncertainty will be maximal ($u_y^x = 1$). Values of $a_y^x > 0.5$ will result in more uncertainty being converted to belief, and vice versa.

5.1.4 Reputation.

Reputation in probabilistic trust systems is calculated by aggregating the evidence in form of $\langle r_y^x, s_y^x \rangle$ from trustful providers [3], [17]. However, since we analyze the effects of disconnection in those initial situations where no prior evidence is available, we also assume that no *reputational* evidence exists anywhere within the society. Moreover, it is pretty obvious that aggregating the evidence provided by the unknown agents in the group may lead to weak or unreliable reputations. Hence, for peer selection, we only use the local trust values computed at each trustor.

5.2 Stereotype Function

To incorporate our estimates into the trust bootstrapping process, we use the base rate as in [8]. For a given trustee y, the base rate $a_y^x = f(\vec{A}_y)$. When no evidence has been accrued for trustee y we have maximum ambiguity, i.e. $w_y^x = \langle 0, 0, 1, 0.5 \rangle$. In this case, a_y^x alone determines the value of $P\left(w_y^x\right)$. However, as more evidence is obtained, the value of u_y^x decreases, and so does the weight carried by a_y^x in the trust value. This fulfills the requirement that the effect of our estimates diminishes as direct evidence is accrued. We refer to this condition as Requirement 1.

Another central observation to be made is that activity stereotypes are useful to form a tentative estimate of the '*No Response*' component of trust evaluation. This means that the increase of the base rate above its default value a_{def} must never preclude the estimated trust value from reducing quickly if the agent misbehaves just in the first interaction. Otherwise, a malicious agent can perform an activity where agents have historically stayed connected for long time to attract trustors, and then behave badly.

Although such a requirement is inherently subsumed by Requirement 1, we have derived an upper bound on the base rate a_y^x that ensures that estimated trust values rapidly fall below the default trust value a_{def}. That is, in the absence of evidence and stereotypical knowledge, the default trust value for any unknown agent in the society is $P\left(w_y^x\right) = b_y^x + a_{def} \cdot u_y^x = a_{def}$, since $b_y^x = 0$ and $u_y^x = 1$.

The upper bound ensures that a stereotypical estimate for $P\left(w_y^x\right)$ drops below a_{def} after just I consecutive bad transactions. The value of I can be seen as an expression of the misprediction risk. A small I signals that malicious trustees can be rapidly discarded. We take as a reference the value of a_{def} because stereotypes never decrease the value of a_y^x below a_{def}. This is explained in detail in the next section. We term this requirement "Requirement 2".

Lemma 1. Given default trust value $a_{def} \in [0, 1]$, fulfilling Requirement 2 requires the base rate a_y^x to be:

$$a_y^x \leq \min\left\{ a_{def} \frac{(I+2)}{2}, 1 \right\}. \tag{7}$$

Further, there exists an upper bound I_{max} on the number of negative interactions I beyond which Requirement 2 is always satisfied: $I_{max} = \frac{2}{a_{def}} - 2$.

Proof: A rigorous proof is given in Appendix D. □

As in many trust systems the default trust value a_{def} is usually 0.5 to keep neutrality, it is interesting to know how many interactions are needed to drop $P\left(w_y^x\right)$ below 0.5. From (7), it follows that $I = 2$ consecutive negative interactions are sufficient, even if stereotypical prediction raised $P\left(w_y^x\right) = a_y^x = f(\vec{A}_y)$ from 0.5 to 1. This certifies that predictions on the '*No Response*' component of trust do not affect the trust formation process in initial cases.

5.2.1 Stereotypical Estimation

Based on the above observations, we are ready to discuss on the shape of the stereotypical function f. The function takes as input a vector of activities $\vec{A}_y$ that are currently being executed by an agent y and returns a prediction of its susceptibility to participate continually.

The computation of the predicted value is as follows. For each activity A_i in $\vec{A}_y$, the trustor first calculates the probability that an agent conducting activity A_i remains connected for a duration of ℓ transactions. Formally, the time for the ℓ^{th} arrival is distributed as an Erlang-ℓ and has mean $\frac{\ell}{\lambda}$. For simplicity, we will use the mean as an input time to compute this probability, denoted by $p_i^{\ell/\lambda}$.

TABLE 2
Activity profile.

Activity	r_{on}	Mean ontime	r_{off}	Mean offtime
A_1	0.2	5 hours	1	1 hour
A_2	4	15 minutes	2	30 minutes

We consider that this value is obtained by asking one of the trusted parties who keep a record of the session durations of the prior agents that performed the activity. Notice that from the uptime session durations of agents, probability $p_i^{\ell/\lambda}$ can be easily obtained by computing:

$$F^c_{A_i}(\ell/\lambda) = \Pr\left(L > \ell/\lambda \middle| A_i\right),$$

where $F^c_{A_i}$ denotes the empirical complementary uptime distribution for the agents who participated in activity A_i in the past. For the acquisition of uptime session lengths, we assume the existence of a secure monitoring protocol for agent availability like AVMON [26] augmented with proof of interactions [27]. These signed certificates can be presented to prove engagement in a specific activity.

Next, the maximum of the probabilities $p_i^{\ell/\lambda}$ is picked and normalized to fit the real interval $[a_{def}, a_{max}]$, where a_{max} is the threshold on the base rate calculated from (7). We chose to use the maximum because it is reasonable to expect that an agent who is involved in several activities at the same time stays connected until completion of the longest activity.

The reason of the normalization is to avoid favoring the new agents for which no activity is known, as these agents are assigned the default base rate a_{def}. As a result, agents who show participation in classified activities will always be preferred over agents for which no activity is known. This should encourage newcomers to participate in known activities. Putting all pieces together, f is given by:

$$f(\vec{A}_y) = \max\left\{p_i^{\ell/\lambda}\right\}_{A_i \in \vec{A}_y} (a_{max} - a_{def}) + a_{def}. \quad (8)$$

As an example, consider a P2P file sharing application and a agent who is performing two activities: A_1 and A_2. A_1 consists of downloading an MP3 file of a famous song while A_2 consists of downloading a large video file. After asking experienced trustors for stereotypical predictions, we get that $p_1^{\ell/\lambda} = 0.01$ while $p_2^{\ell/\lambda} = 0.8$, since activity A_2 requires, in general, more time to accomplish. Assuming $a_{max} = 0.75$ and $a_{def} = 0.5$ (totally ignorant opinion), the predicted $P\left(w_y^x\right)$ is $f(\vec{A}_y) = p_2^{\ell/\lambda}(0.75 - 0.5) + 0.5 = 0.7$.

Finally, it is worth noting that activity stereotypes are little intrusive to trust management. Clearly, stereotyping does not affect the management of the trustworthiness of relationships among agents. The underlying trust system can operate without them. Stereotypes only reshape trust evaluation of unknown agents to avoid experiencing too many disconnections during trust bootstrapping. Only in the particular situation that stereotypical opinions could be communicated within the society, a more specific trust management will be needed, as it happens with classical stereotypes [8].

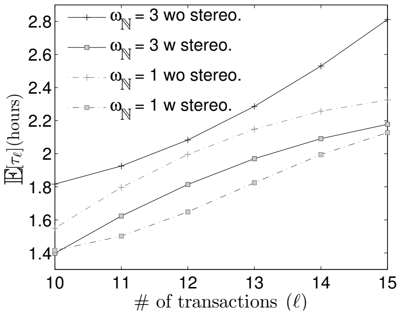

Fig. 3. Mean bootstrapping time $\mathbb{E}\left[\tau_\ell\right]$ as ℓ is varied, with and without stereotypes, for two distinct penalties for '*No Response*': $\omega_{\mathbb{N}} = 1$ (moderate) and $\omega_{\mathbb{N}} = 3$ (aggressive).

TABLE 3
Simulation parameters.

Parameter	Value	Description
N_{groups}	250	Ad-hoc group count
G_{size}	10	Ad-hoc group size
$p_{\mathbb{B}}$	0.5	Fraction of malicious trustees
λ	10	Transactions per hour
a_{def}	0.5	Default trust value
a_{max}	0.75	Threshold on base rate
$\omega_{\mathbb{B}}$	1	Penalty for bad response

6 EVALUATION

6.1 Experimental Setup

We validated our approach with exhaustive simulations. Here we only report a subset of the results. It is worth to note here that because we have been the first to tackle the problem of disconnection during trust bootstrapping, we found no way of *fairly* benchmarking activity stereotypes against the existing literature. A clarifying discussion on this issue is given in Appendix G.

In our experiments, we simulated a system composed of $N_{groups} = 250$ groups. In each group, a trustor wanted to receive service from G_{size} agents from whom no direct and reputational evidence was forthcoming. The goal of the trustor was to produce accurate trust evaluations.

6.1.1 *Activity Profiles and Threat Model*

Each of the G_{size} agents in each group was assigned an activity profile which specified how long it will be online and disconnected. Each profile specified two parameters: the parameters of the two exponential distributions from which online and offline durations were drawn. Activies profiles can be composed of none, one, or two activities. The test profiles used in our experiments are reported in Table 2. As reflected in the table, agents who historically carried out activity A_1 exhibited ontime durations drawn from CDF $1 - e^{-0.2x}$ and disconnection durations drawn from CDF $1 - e^{-x}$, which corresponds to a mean ontime and offtime of 5 and 1 hours, respectively. Conversely, the agents who historically participated in activity A_2 in the past presented a high turnover rate with an average ontime and offtime of 15 and 30 minutes, respectively. To make the identification of the agents with longer ontimes

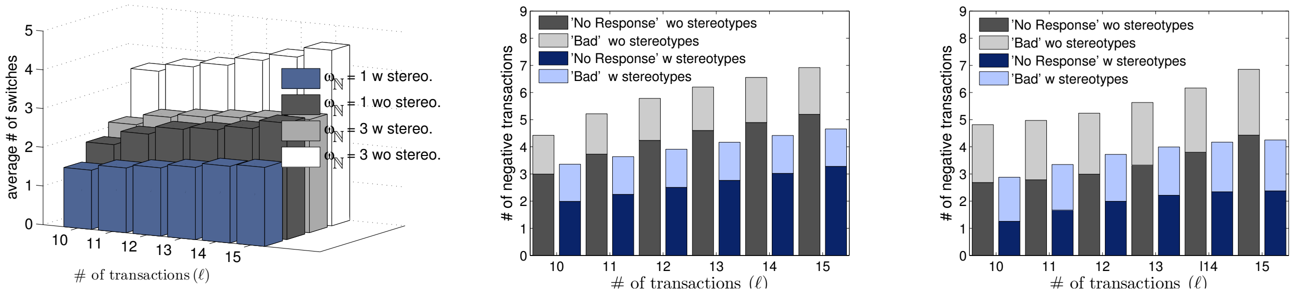

(a) Mean number of switches (b) For $\omega_{\mathbb{N}} = 1$. (c) For $\omega_{\mathbb{N}} = 3$.

Fig. 4. Effectiveness as ℓ is varied, with and without stereotypes, and for two distinct penalties for '*No Response*': $\omega_{\mathbb{N}} = 1$ (moderate) and $\omega_{\mathbb{N}} = 3$ (aggressive). (a) Mean number of switches. (b)-(c) Mean number of negative transactions.

harder, only 20% of the trustees within each group were assigned activity A_1. In practice, this represented that the probability of bootstrapping trust with little interruption was at most of 20% in the absence of stereoytpes. Such a of activities is enough to approximate real user behavior in P2P applications such as BitTorrent, where a swarm is nothing but a group of unknown users, or in file-sharing systems such as Gnutella and Kazaa [28], to name a few.

Regarding the threat model, while the list of potential attacks is myriad, we use the simplest form of threat that can be identified in any trust system: *a group of malicious agents who always give fraudulent service*, or with sufficient degradation to be qualified as a negative interaction. The main reason for using such a simple model was to isolate any effect caused by the trust model from stereotyping, whose aim is not to predict trustworthiness but minimize disconnection.

Specifically, a fraction $p_{\mathbb{B}}$ of the agents were disposed to act maliciously in each simulation run. By default, we set $p_{\mathbb{B}}$ to 0.5 in order to have half of the agents in activity A_1 provide bad service and evaluate if Requirement 2 is fulfilled. If true, this implies that our approach is able to react quickly against the trustees who were predicted to stay connected but render poor service, that is, when *no activity-behavioral correlations exist*.

6.1.2 Trust Model

Before interacting with any agent in a group, uncertainty on the trustworthiness of each agent is maximal. Hence, the opinion held by a trustor x about each agent y in the group is $w_y^x = \langle 0, 0, 1, a_y^x \rangle$. When stereotypes are applied, the base rate $a_y^x = f(\vec{A}_y)$, which may favorably bias trust evaluation $P\left(w_y^x\right)$. When stereotypical prediction is not possible, a_y^x is simply equal to a_{def}. In our experiments, we set a_{def} to the uninformative prior 0.5, which meant that before any interaction took place, both positive and negative outcomes were considered equally like.

According to Lemma 1, we fixed a_{max} to 0.75 to ensure that stereotypical misprediction could be corrected if the result of the first interaction with an unknown agent was negative.

By default, the penalty for bad response $\omega_{\mathbb{B}}$ was set to 1, to equalize it with positive interaction. Observe that in the underlying trust model, the degree of belief, disbelief and uncertainty that constitute an opinion are calculated by operating directly on the number of observed positive and negative experiences, which implicitly assumes that the magnitude of reward obtained and penalty incurred is of 1. Values of $\omega_{\mathbb{B}} > 1$ will result in greater punishment for bad response compared with the reward received for positive behavior. Default parameters for simulations are summarized in Table 3.

Last but not least, we assume that the trustor in each group selects the most trusted agent at every interaction, i.e., the agent with the highest trust value $P\left(w_y^x\right)$, which is the most common decision model in trust literature.

6.2 Results

Here we present the result of our experiments. The main hypothesis to validate is *whether trust bootstrapping will be better with stereotypes than without stereotypical information*. In this sense, there are two important aspects to evaluate. On one hand, if *stereotypical* biases are present, then trust bootstrapping times should be shorter, and on the other hand, if stereotypes reduce the number of unsatisfactory interactions.

Due to space constraints, we report here a small subset of the compiled results to give a sense of the advantages of activity stereotypes. The rest of the results are given in Appendices H-J. Appendix H examines the effect of the fraction of malicious agents $p_{\mathbb{B}}$. Appendix I studies what happens when the penalty $\omega_{\mathbb{B}}$ is greater than $\omega_{\mathbb{N}}$. Finally, Appendix J shows the time evolution of no responses.

6.2.1 Experiment I: General Effectiveness

Fig. 3 illustrates the average bootstrapping time for two values of the disconnection penalty: $\omega_{\mathbb{N}} = 1$ and $\omega_{\mathbb{N}} = 3$, which represent a normal and an aggressive punishment, respectively. Recall that the penalty for malicious service $\omega_{\mathbb{B}}$ was fixed to 1. Therefore, a disconnection penalty $3X$ greater than $\omega_{\mathbb{B}}$ can be considered aggressive. Besides the obvious conclusion that informed trustee selection based on activity stereotypes performs significantly better, one important observation should be made about this result. Such an observation is the *empirical evidence* that contrary to intuition but consistent with our analysis in Section 2, *an aggressive punishment can increase the trust bootstrapping time*, instead of reducing it. As in the figure, the average bootstrapping time is longer for $\omega_{\mathbb{N}} = 3$ than for $\omega_{\mathbb{N}} = 1$. The reason is that the trustees present a marked tendency to return sooner rather than later in both activities, which favored waiting for a trustee to come back online in front of switching to a new trustee. Recall that a higher level of punishment means a larger number of switches due to stronger drops in trust values. This is further evidenced in Fig.4(a), where the number of agent switches increases with the intensity of the punishment. More importantly, activity stereotypes reduce the number of agent switches caused by disconnections, which explains the reduction in trust bootstrapping times.

Another interesting indicator is the average number of negative transactions experienced by a trustor during the trust bootstrapping process. The added value of this indicator is that gives us a clear hint about the effect that disconnection punishment has on service continuity.

Fig. 4 illustrates this effect and gives a clear picture of the existing *tradeoff* between service continuity and risk. This tradeoff is easy to observe by visual comparison of Fig. 4(b) with Fig. 4(c). Fig. 4(b) plots the average number of negative transactions for moderate punishment across all groups. Fig. 4(c) does so for aggressive punishment. By comparing Fig. 4(b) with Fig. 4(c), it is easy to observe that although a higher penalty reduces the occurrence of service interruptions, it unfortunately heightens the risk of receiving a wrong response, and vice versa for lower penalties. Notice that while in Fig. 4(b) the portion of the bars for bad responses is comparatively smaller than in Fig. 4(c), the inverse result is obtained for no responses. Either way, *activity stereotypes are able to reduce the presence of service interruptions*, making more attractive moderate punishment because of the smaller risk of bad interaction it entails.

Overall, our results show that *activity stereotyping offers a clear improvement in the initial cases where prior direct and reputational evidence is lacking*, and we believe it is a line of work to be further explored.

7 Conclusions

In this work, we have analyzed to what extent punishing disconnection affects trust formation in these initial cases where prior evidence is not available, an issue that has been overlooked in the literature. First, we have proven analytically the lack of a *universally optimal disconnection penalty* and shown its dependence on the connection and disconnection habits of users. Second, we have presented a new mechanism based on activity stereotypes to make trust bootstrapping *quicker* and *less dependent* on the way trust systems punish disconnection.

Acknowledgements

This research work was partly sponsored by the EC FP7 through project CloudSpaces ($\text{FP7}-317555$). The authors would like to thank the anonymous reviewers for their comments, which helped to improve the article greatly.

References

[1] D. P. Anderson, “Boinc: A system for public-resource computing and storage,” in *GRID*, 2004, pp. 4–10.

[2] W. Wang et al., “Ripple-stream: Safeguarding p2p streaming against dos attacks,” in *ICME*, 2006, pp. 1417–1420.

[3] Z. Liang and W. Shi, “PET: A PErsonalized Trust Model with Reputation and Risk Evaluation for P2P Resource Sharing,” in *HICSS*, 2005, pp. 201b–201b.

[4] N. Fedotova and L. Veltri, “Reputation management algorithms for dht-based peer-to-peer environment,” *Comput. Commun.*, vol. 32, no. 12, pp. 1400–1409, 2009.

[5] X. Li et al., “A multi-dimensional trust evaluation model for large-scale p2p computing,” *J. Parallel Distrib. Compu.*, vol. 71, no. 6, pp. 837–847, 2011.

[6] D. Stutzbach and R. Rejaie, “Understanding churn in peer-to-peer networks,” in *IMC*, 2006, pp. 189–202.

[7] X. Liu, A. Datta, K. Rzadca, and E.-P. Lim, “Stereotrust: a group based personalized trust model,” in *CIKM*, 2009, pp. 7–16.

[8] C. Burnett, T. J. Norman, and K. Sycara, “Bootstrapping trust evaluations through stereotypes,” in *AAMAS*, 2010, pp. 241–248.

[9] Z. Liang and W. Shi, “Analysis of ratings on trust inference in open environments,” *Perform. Eval.*, vol. 65, no. 2, pp. 99–128, 2008.

[10] M. Sánchez-Artigas, P. García-López, and B. Herrera, “Exploring the feasibility of reputation systems under churn,” *IEEE Comm. Letters*, vol. 13, no. 9, pp. 558–560, 2009.

[11] M. Sánchez-Artigas, “Churn delays uncertainty decay in p2p reputation systems,” *IEEE Internet Computing*, vol. 99, 2010.

[12] M. Sánchez-Artigas and B. Herrera, “Understanding the effects of p2p dynamics on trust bootstrapping,” *Information Sciences*, vol. 236, no. 0, pp. 33–55, 2013.

[13] R. Hermoso, H. Billhardt, and S. Ossowski, “Role evolution in open multi-agent systems as an information source for trust,” in *AAMAS*, 2010, pp. 217–224.

[14] R. Conte and M. Paolucci, *Reputation in Artificial Societies: Social Beliefs for Social Order*. Kluwer Academic Publishers, 2002.

[15] C. Burnett, T. J. Norman, and K. Sycara, “Sources of stereotypical trust in multi-agent systems,” in *TrustCom*, 2011, pp. 25–39.

[16] M. Sensoy, B. Yilmaz, and T. J. Norman, “Discovering frequent patterns to bootstrap trust,” in *ADMI*, 2012, pp. 93–104.

[17] A. Jøsang, R. Hayward, and S. Pope, “Trust network analysis with subjective logic,” in *ACSC*, 2006, pp. 85–94.

[18] Z. Yao et al., “Modeling heterogeneous user churn and local resilience of unstructured p2p networks,” in *ICNP*, 2006, pp. 32–41.

[19] A. Datta and K. Aberer, “Internet-scale storage systems under churn – a study of the steady-state using markov models,” in *P2P*, 2006, pp. 133–144.

[20] R. Kumar, Y. Liu, and K. Ross, “Stochastic fluid theory for p2p streaming systems,” in *INFOCOM*, 2007, pp. 919–927.

[21] Z. Yang et al., “Exploring peer heterogeneity: Towards understanding and application,” in *P2P*, 2011, pp. 20–29.

[22] J. Mundinger and J.-Y. L. Boudec, “Analysis of a reputation system for mobile ad-hoc networks with liars,” *Perform. Eval.*, vol. 65, no. 3-4, pp. 212–226, 2008.

[23] C. M. Jonker and J. Treur, “Formal analysis of models for the dynamics of trust based on experiences,” in *MAAMAW*, 1999, pp. 221–231.

[24] G. Song et al., “Replica placement algorithm for highly available peer-to-peer storage systems,” in *AP2PS*, 2009, pp. 160–167.

[25] S. Marti and H. Garcia-Molina, “Taxonomy of trust: categorizing p2p reputation systems,” *Comput. Netw.*, vol. 50, no. 4, pp. 472–484, 2006.

[26] R. Morales and I. Gupta, “Avmon: Optimal and scalable discovery of consistent availability monitoring overlays for distributed systems,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 20, no. 4, pp. 446–459, 2009.

[27] A. Singh and L. Liu, “Trustme: Anonymous management of trust relationships in decentralized p2p systems,” in *P2P*, 2003, pp. 142–149.

[28] K. P. Gummadi et al., “Measurement, modeling, and analysis of a peer-to-peer file-sharing workload,” *SIGOPS Oper. Syst. Rev.*, vol. 37, pp. 314–329, 2003.

PLACE
PHOTO
HERE

Marc Sánchez-Artigas received the PhD degree in 2009 from the Universitat Pompeu Fabra, Barcelona, Spain. During his PhD, he worked at École Polytechnique Fédérale de Lausanne (EPFL) under the supervision of professor Karl Aberer. In 2009, he joined the Universitat Rovira i Virgili as an assistant professor. He received the Best Paper Award from the 32th IEEE Conference of Local Computer Networks (LCN) held in Dublin, 2007. He also served as a Guest Editor for Computer Networks Journal, Elsevier. He organized the 12th edition of the IEEE P2P 2012 conference that was held at Tarragona, Spain. He is actively involved in the coordination of the European FP7 project called CloudSpaces on personal Clouds.

PLACE
PHOTO
HERE

Blas Herrera received the PhD degree in differential geometry from the Universitat Autónoma de Barcelona, Spain, in 1994. He is currently an active researcher on Mechanics, Geometry and Computer Engineering.