

Securing P2P Systems from Sybil Attacks through Adaptive Identity Management

Weverton Luis da Costa Cordeiro, Flávio Roberto Santos,
Gustavo Huff Mauch, Marinho Pilla Barcelos, Luciano Paschoal Gaspary

Institute of Informatics – Federal University of Rio Grande do Sul
Av. Bento Gonçalves, 9500 – 91.501-970 – Porto Alegre, RS, Brazil
{weverton.cordeiro, frsantos, ghmauch, marinho, paschoal}@inf.ufrgs.br

Abstract—An effective approach to tackle Sybil attacks consists in establishing computational puzzles to be solved prior to granting new identities. Solutions based on this approach, despite their potential, do not distinguish between identity requests originated from correct users and attackers, requiring both to pay the same cost for an identity requested. Assuming computational puzzles of similar complexity, attackers having access to high performance computing hardware might be able to solve them orders of magnitude faster than legitimate users. Consequently, attackers may obtain a larger number of identities. However, simply increasing the complexity of puzzles would hamper the admission of legitimate peers to the network. To address this problem, we propose the use of adaptive computational puzzles as an approach to limit the spread of Sybils. The key idea is to estimate a trust score of the source from which identity requests depart, calculated as a proportion to the recurrence rate of identity requests originated from other sources. The higher the frequency (the) user(s) associated to a source perform(s) identity requests, the lower the trust score of that source and, consequently, the higher the complexity of the puzzle to be solved. Results achieved by means of an experimental evaluation show the effectiveness of our solution. While comparatively more complex puzzles are assigned to potential attackers, legitimate users are minimally penalized with easier-to-solve puzzles.

I. INTRODUCTION

The Sybil attack [1] represents one of the most elementary threats to identity management schemes in P2P networks, and consists in the creation of multiple, counterfeit identities, termed as Sybil identities (or peers). The main motivation for a malicious user to launch such an attack is to control the majority (or a large fraction) of the peer identities in the network. As a consequence, there is a high probability that interactions between legitimate peers will be mediated by any of the counterfeit identities, and manipulated for the benefit of the attacker herself [2]. An attacker controlling several counterfeit identities may also subvert polling-based algorithms, thus manipulating the reputation of peers or contents being shared in the network. In addition, the Sybil attack serves as basis for launching other attacks in P2P networks, such as Eclipse [3], Free-riding [4], and Pollution [5].

A promising approach to protect P2P networks from Sybil attacks consists in assigning computational puzzles to requesting users prior to granting them or renewing their identities [6]. The idea behind the use of computational puzzles is that legitimate peers are able to prove their good intentions with the network, by compromising a fraction of their resources. In contrast, malicious peers interested in creating multiple identities are forced to spend a large fraction of their time

processing puzzles and, therefore, consuming resources. This reduces their power to assume large number of identities.

The idea of using computational puzzles for identity management in P2P networks is not new [6], [7]. Despite their potential, existing proposals do not distinguish between requests from legitimate users and those originated by attackers. Since both are subject to the payment of the same (computational) price for each requested identity, these proposals may lose effectiveness when the computational resources of attackers outweigh those of legitimate users. If computational puzzles are of similar complexity, an attacker with access to high performance computing hardware might be able to solve them orders of magnitude faster than legitimate users. Consequently, an attacker may obtain a larger number of identities. However, simply increasing the complexity of puzzles would hamper the admission of legitimate peers to the network.

In this paper we propose the use of adaptive puzzles as a strategy to limit the spread of Sybils. In contrast to previous approaches, our solution estimates a trust score of the source from which identity requests originate, calculated as a proportion to the recurrence rate of identity requests from other sources. The higher the frequency a source requests identities¹, the lower its trust score and, consequently, the higher the complexity of the puzzle to be solved by (the) user(s) associated to that source. Results achieved by means of a set of experiments – using real traces of identity requests from popular P2P communities – evidence the effectiveness of our solution. While comparatively more complex puzzles are assigned to potential attackers, legitimate users are minimally penalized with easier-to-solve puzzles.

The remainder of this paper is organized as follows. Section II briefly reviews related work that are closer to this paper. Section III presents the proposed solution for adaptive puzzles as a protection against Sybil attacks, whereas Section IV describes the evaluation carried out to evaluate its effectiveness. Finally, Section V closes the paper with concluding remarks and prospective directions for future research.

¹In the context of this work, “a *source* requests identities” means in fact “user(s), from a certain *source*, request(s) identities”. *Source* may refer to a user’s workstation, a local network, an Autonomous System (AS), etc. (identified by an IP address or prefix). In substitution or as a complement, *source* may be associated with a network coordinate provided by a system such as Vivaldi [8] and Veracity [9].

II. RELATED WORK

Investigations carried out to tackle the Sybil attack may be classified according to the strategy employed to enforce peer authenticity. The main categories correspond to solutions based on *weak* and on *strong identities*. The first category comprises solutions in which peers have autonomy to create their own identities. In this case, there is an estimate of the number of Sybil identities acceptable. An estimate may be useful, for instance, for applications that can tolerate a certain fraction of Sybil nodes. Examples in this category include those proposed by Yu *et al.* [10] and by Danezis *et al.* [11]. One of these mechanisms, SybilLimit [10], explores the concept of social networks to limit the spread of Sybil identities in the P2P community, also estimating an upper bound for the number of counterfeit identities that are “accepted”.

In the second category of solutions, peers may only obtain identities through Certification Authorities. Its main advantage is the difficulty imposed to peers that attempt to create and control several identities, or take control of someone else’s identity. Nevertheless, such solutions may limit the scalability of the P2P network, force users to trust unknown certification authorities, and hamper the access of potential users (for example, when she needs to provide personal data or to pay fees to obtain one identity). The proposals that fit in this category attempt to minimize some of these side-effects. For example, in [12] and [13], the authors have focused on the decentralization of the Public-Key Infrastructure (PKI). Still, these proposals require the exchange of a high amount of messages between peers, and rely on the contribution of a minimal number of peers to operate as expected.

In the paper that describes the Sybil attack [1], Douceur demonstrated that a robust and scalable solution for peer authentication in P2P networks requires a certain degree of centralization (for example, by employing certification authorities). Since then, and considering the severe constraints imposed by the use of certification authorities, an intermediate category of solutions that has attracted attention in recent years is to condition identity granting/renewal to the previous resolution of computational puzzles. Puzzles are typically assigned to users by a *bootstrap service*, and aim at decreasing malicious users’ capabilities of creating counterfeit identities without compromising the intrinsic characteristics of P2P networks (such as scalability, decentralization, and peer autonomy).

Approaches based on computational puzzles have presented satisfactory results when using challenges created and/or verified in a distributed fashion. Borisov [6], for example, has showed the technical feasibility of using puzzles generated periodically and distributedly, by proposing an approach in which peers that participate in the puzzle generation process are able to validate its solution. Rowaihy *et al.* [7], in turn, have proposed an approach based on multiple puzzle generation entities. It requires that, prior to obtaining identities, users contact one of these entities and solve a series of puzzles.

Nonetheless, existing solutions do not address the issue of adjusting puzzle complexity. More specifically, when using puzzles of equal computational complexity to all users, it becomes very hard to choose a complexity that is adequate to the entire system. On one extreme, puzzles having higher complexity penalize legitimate users with less powerful hardware. On the other extreme, puzzles having lower complexity

may not hamper attackers with high performance computing hardware. Therefore, the use of a uniform complexity for puzzles may benefit attackers at the expense of legitimate users. The major contribution of this paper is to parameterize the complexity of puzzles according to the behavior of each source in the network. Users associated to sources whose behavior is more similar to the average behavior of other sources are benefited with less complex puzzles. In contrast, users associated to sources whose behavior deviates significantly from the behavior of other sources are forced to cope with more complex puzzles to obtain identities.

III. PROPOSED SOLUTION

Given the problem of the Sybil attack in P2P networks, the solution proposed in this paper is to establish adaptive computational puzzles for identity management. It involves solving three main sub-problems, namely: (i) characterize the behavior of sources (or the behavior of users associated to them); (ii) calculate the trust score of a source based on the observed behaviors; and (iii) deal with the dynamics of users’ behavior in the calculation of the trust score. Each of these sub-problems are addressed in the following subsections.

A. Employing Recurrence Rates to Characterize Behaviors

In order to enable the characterization of sources of identity requests, two metrics are introduced: *source recurrence rate* (ϕ) and *network recurrence rate* (Φ). The former represents the frequency in which identity requests to the P2P bootstrap service originate from some specific source, averaged within a given time interval t_w (with $t_w > 0$). The latter corresponds to the average frequency in which all sources make use of the bootstrap service to request new identities.

Network recurrence rate is computed following Equation 1, which employs the harmonic mean of the recurrence rates of all sources. In the equation, ϕ_i represents the recurrence rate of the i -th source in the P2P network. Note that when $\phi_i = 0$ users associated to source i have not requested any identity; such a source can be safely ignored.

$$\Phi = \frac{n}{\sum_{i=1}^n \frac{1}{\phi_i}} \quad (1)$$

B. Calculating Trust Scores from Observed Behaviors

For an attacker to succeed in a Sybil attack, she must request a large number of identities to the bootstrap service. This behavior leads to the increase in the recurrence rate observed for the source associated to the attacker. Conversely, it is expected that the sources associated to legitimate users perform fewer identity requests (for example, in the moment these users register themselves in the P2P network). Therefore, the underlying idea to control Sybil attacks is to assign more complex puzzles to users associated to sources whose recurrence rates are higher than the network recurrence rate.

By comparing the behavior of each source (inferred from ϕ) and the presumed normal network behavior (inferred from Φ), we calculate the *relationship between source and network recurrence rates* (ρ). Obtained from Equation 2, it assumes values lower than zero to denote how many times the recurrence rate of the i -th source is lower than the network recurrence

rate. Likewise, ρ higher than zero expresses how many times higher is the recurrence rate of the source.

$$\rho = \begin{cases} -\frac{\Phi(t)}{\phi_i(t)} & \text{if } \phi_i(t) < \Phi(t) \\ \frac{\phi_i(t)}{\Phi(t)} & \text{if } \phi_i(t) \geq \Phi(t) \end{cases} \quad (2)$$

The relationship between the source and network recurrence rates (ρ) is used to compute the *source trust score* (θ). This score, calculated according to Equation 3, assumes values in the interval $[0, 1]$: on one extreme, values closer to 1 denote a higher trust on the legitimacy of (the) user(s) associated to that source; on the other, values closer to 0 indicate higher distrust, i.e., a higher probability that (the) user(s) associated to that source is(are) launching a Sybil attack. Equation 3 is normalized so that the extreme values, 0 and 1, represent total distrust and total trust upon a given source, respectively.

$$\theta(t) = 0.5 - \frac{\arctan(a \times (\rho - c)^{(1+2 \times b)})}{\pi} \quad (3)$$

Figure 1 presents four different configurations that illustrate how the trust score obtained for a given source varies as a function of ρ . In each of these configurations, the terms a , b , and c of Equation 3 assume arbitrary values and play an important role in controlling how fast the curve decreases, its amplitude and translation, respectively.

The plots in Figure 1 reveal two important properties that Equation 3 holds. The first – and most important – property corresponds to the fact that curves are asymptotic in 0 and 1. Therefore, for $\rho \rightarrow -\infty$ or $\rho \rightarrow +\infty$, there is always a corresponding value for the trust score. The second one refers to the minimal variations of the trust score for values of ρ closer or equal to 0, situation in which the source behaves similarly or equals to the network average. This property allows a certain degree of tolerance in the evaluation of source behavior. To illustrate, consider the configuration ($a = 0.001$, $b = 2$, $c = 0$) shown in Figure 1; variations of ρ within the interval $-2 \leq \rho \leq 2$ have minimal impact, since these are slightly similar to the pattern observed in the network. Behaviors that deviate significantly from this interval, however, will be assigned lower (or higher) values for the trust scores. This is also the case, for example, of the abrupt variations in configuration ($a = 0.001$, $b = 2$, $c = 0$) within the intervals $-5 \leq \rho \leq -2$ and $2 \leq \rho \leq 5$.

C. Dealing with the Dynamics of Users' Behavior

Peer autonomy is an important characteristic of P2P networks. It means peers may join and leave the network at any moment, on purpose, or become unavailable. It also means peers may not rely on external entities to make any decisions. One possible effect of such dynamics is the variation in the behavior pattern of both individual sources and the network as a whole. Next, we discuss how the proposed solution deals with the dynamics of observed behaviors.

Although Equation 3 allows us to establish the trust score of a given source at instant t , it does not take into account the source historical behavior. Hence, a smoothing factor is used to properly represent the trust score of a given source in light of its historical behavior. The smoothed trust scores is calculated as shown in Equation 4. The smoothing factor β determines the weight of past behavior in the calculation of

the trust score at the present instant (t), assuming values in the interval $(0, 1]$. On one extreme, values of β closer to 0 assign a higher weight to the historical behavior of the source under consideration. On the other extreme, values of β closer to 1 assign a higher weight to the current behavior of the source. In the special case in which $\beta = 1$, the current trust score (as calculated through Equation 3) is fully considered, and the historical behavior, totally ignored.

$$\theta_s(t) = \beta \times \theta(t) + (1 - \beta) \times \theta(t - 1) \quad (4)$$

Another important issue, still concerning the dynamics of network behavior, lies in the fact that recurrence rates may vary across different times of the day, month or year. In order to address seasonality in the pattern of identity requests, a sliding window – a time interval which starts t_w hours in the past – is used to restrict the amount of requests considered in the calculation of the recurrence rate ϕ of each source and, consequently, the network recurrence rate, Φ (as discussed in Subsection III-A). The window slides forward in time steps of duration t_d (with $t_d \leq t_w$). Older requests are gradually discarded, allowing room to newer ones, which are more representative of the current state of the P2P network.

IV. EVALUATION OF THE PROPOSED SOLUTION

In order to evaluate the technical feasibility of using adaptive puzzles to tackle Sybil attacks in P2P networks, we have developed a prototypical implementation of a bootstrap service. In summary, the developed service aggregates the functionalities of management of identity requests from users interested in joining the network, assignment of puzzles for each identity request, validation of solutions received for assigned puzzles, and granting (or denial) of requests (according to the correctness of received solutions).

An instance of the bootstrap service has then been used to carry out several experiments, using a combination of both realistic traces of identity requests and synthetic ones. The experiments aimed at confirming that (i) puzzles proposed to legitimate users minimally penalize them; (ii) puzzles assigned to potential attackers have comparatively higher computational complexity; and (iii) the proposed solution is robust and resilient even when there is a large fraction of attackers in the system.

The remainder of this section is organized as follows. Subsection IV-A describes the details of the environment considered in the analysis (characteristics of historical traces of identity requests employed, parameter setting, etc.). Subsection IV-B, in turn, presents and discusses the results obtained with the proposed solution.

A. Configuration of the Experimental Environment

Table I presents a summary of the characteristics of the trace employed in the experiments, values for each of the parameters involved in the solution, and the characteristics of the Sybil attacks considered in the analysis. Each of these aspects is discussed in detail in the following paragraphs.

The experiments were performed based on historical traces of identity requests obtained from the closed P2P community Bitsoup [14]. Since the admission (creation of new accounts) in this community is moderated, we assumed that the trace under analysis did not contain any records of Sybil attacks.

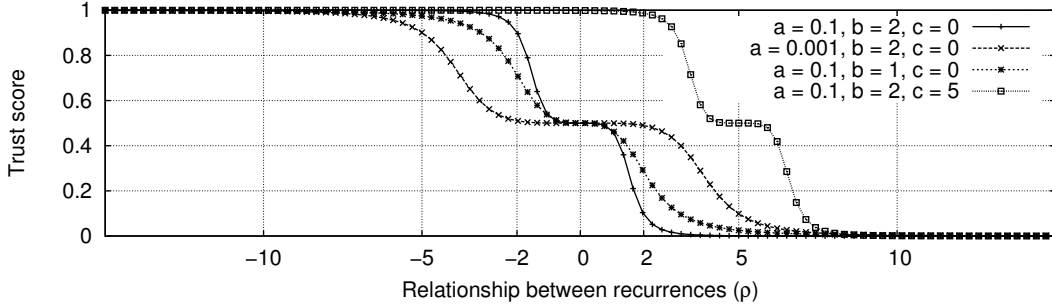


Fig. 1. Sample values for parameters a , b , and c of Equation 3 for the calculation of the source trust score

TABLE I
INFORMATION REGARDING THE ENVIRONMENT CONSIDERED IN THE
EXPERIMENTAL EVALUATION.

Characteristics of the Trace Employed			
Duration	15 days		
Amount of identities requested	625,079 identities		
Number of distinct sources	44,315 sources		
Mean time between requests	2.08 seconds		
Average number of identity requests	14.21 identities		
Parameters of the Proposed Solution			
a	0.1	b	2
c	5	β	0.125
Window Duration (t_w)	8 hours	Step Duration (t_d)	1 hour
Attack Strategies			
Request Rate per Attack Source	1; 1.25; 1.5; 2; and 2.5 requests/hour		
Number of Attack Sources	1; 100; 500; 1,000; and 2,000		

This assumption relies on the notion that the cost necessary to create and manage several counterfeit identities in a closed and moderated file sharing community is prohibitive and, therefore, unlikely to happen. Please observe that such a trace is used solely as an input for emulating the behavior of legitimate P2P users, whose identity requests should never be blocked.

The traces employed in our analysis register activities of identity requests in a time window of 15 days. In this period, 625,079 identities were requested to the bootstrap service, by 44,315 distinct sources; this results in approximately 14 identity requests per source on average, and a global rate of 1 identity being requested every 2.08 seconds.

The parameters involved in the experimental evaluation were defined as follows. The smoothing factor β was defined as 0.125, i.e., the historical behavior of the trust score accounts for 87.5% of the current trust score. This value has showed to be adequate, after successive experiments (omitted in this paper due to space constraints), in order to prevent that sources with historic of misbehavior reach higher values for the trust score when suddenly becoming “well behaved”. The parameters a , b , and c , in turn, were assigned the values 0.1, 2, and 5, respectively, aiming at controlling how the relationship between recurrence rates maps into source trust score. Hence, a value for the source recurrence rate similar or equal to the network recurrence rate ($\phi \simeq \Phi$) results in a value for the trust score closer to 1 (for example, see configuration $a = 0.1$,

$b = 2$, and $c = 5$ in Figure 1). Finally, the sliding window has a duration of 8 hours ($t_w = 8 \times 60 \text{ min}$) and slides hourly ($t_d = 1 \times 60 \text{ min}$). These settings were shown to properly capture the past behavior of each source, whereas disregarding identity requests that no longer represent the current state of the P2P network.

In order to evaluate scenarios in which the P2P network is under Sybil attacks, we injected artificially generated malicious identity requests, considering two distinct strategies. In the first one, an attacker launches a Sybil attack from a single source. In the second strategy, the attacker controls a certain number of sources, and the Sybil attack is launched distributedly; each source requests a small amount of identities, in order to remain unsuspecting in the network. As a result of the evaluation, we observed the amount of identities the attacker is able to obtain *versus* the complexity of the puzzle that the system assigned to requests originated from the sources involved in the attack.

B. Results Obtained and Analysis

In order to organize the discussion of the results obtained, we first observe the values of trust score calculated for legitimate users in the absence of Sybil attack. Subsequently, we evaluate the influence of attacks, originated from a single source, on the trust scores computed to both legitimate users and user(s) associated to that malicious source. Finally, we analyze the resilience of our solution in situations in which the attacker harnesses the power of colluding nodes (e.g., by means of botnets) to obtain a large number of identities.

1) *Overhead Caused to Legitimate Users in the Absence of Sybil Attacks:* Figure 2 shows the complementary cumulative distribution function (CCDF) of the trust scores calculated for identity requests coming from (presumably legitimate) sources of the studied trace. It is important to emphasize that this result refers to identity requests present in the original trace only, and it was not disturbed by occurrences of Sybil attacks.

One may note in the plot shown in Figure 2 that the majority of identity requests achieved high values for the trust score. For example, approximately 45% of the identity requests were performed by users associated to sources having trust score higher or equal to 0.9. This percentage increases to 60% if we consider requests for which a trust score higher or equal to 0.7 is calculated, and to approximately 75% if we consider those with trust scores higher or equals to 0.5. In summary, a significative fraction of sources may obtain computational puzzles of lower complexity, thus causing minimal overhead to users associated to them.

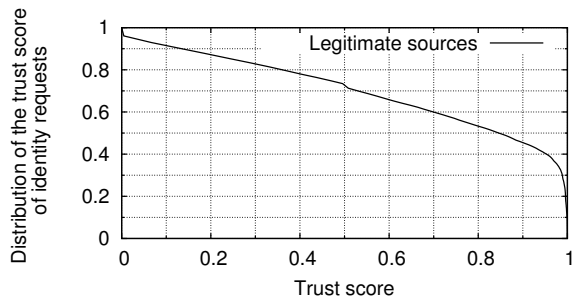


Fig. 2. CCDF of the trust score of requests performed by legitimate sources

2) *Impact to Potential Attackers:* Figure 3 presents the results achieved for five new scenarios, each having an artificially generated Sybil attack. These attacks were launched by a single (malicious) source. The main difference between the scenarios was the identity request rates adopted: 1; 1.25; 1.5; 2; and 2.5 requests per hour, respectively.

An important observation regarding the results presented in Figure 3 corresponds to the influence of the Sybil attack on the trust score achieved by legitimate sources. Regardless of the recurrence rate adopted by the attacker, the distribution of the trust score of legitimate users is always the same. This is due to the resilience of the harmonic mean – statistical measure employed to calculate the network recurrence rate, as discussed in Section III-A – to recurrence rates that deviate significantly in comparison to other sources. For the sake of legibility, only one curve is presented in Figure 3 to illustrate the distribution of the trust score of legitimate sources.

The results presented lead to two distinct conclusions. The first conclusion is that the proposed solution reacts adequately to increases in the sources recurrence rates, severely penalizing those sources that request identities to the bootstrap service in a frequency higher than the network average. The second one is that the solution forces sources to “behave adequately” – i.e., make use of the bootstrap service harmonically in comparison to other sources – in order not to be penalized with more complex computational puzzles.

3) *Resilience of the Proposed Solution to Collusion Attacks:* In this scenario, instead of increasing the recurrence rate to obtain more counterfeit identities, the attacker acts in collusion with other attackers (or makes use of a botnet composed of several zombie stations connected to the Internet). When doing so, the attacker is able to increase the speed in which she obtains identities in the P2P network. Further, she may also alter the “perception of normality” in the network: a higher number of malicious sources behaving similarly in the network tends to change the perception of what is, effectively, the behavior of the majority of sources in the network.

Figure 4 presents the results achieved considering the new strategy of attack. Four scenarios are considered, each having a different number of malicious sources available for the attack: 100; 500; 1,000; and 2,000 sources. In all scenarios, each malicious source requests an average of 1.5 identities per hour. This rate was chosen because it enables the attacker to obtain a significant number of identities whilst not compromising its trust score (as evidenced in the previous analysis).

Observe in Figure 3 that, even using an extremely high

amount of sources, the impact that the attacker causes upon the network normality pattern is relatively limited. For example, in Figure 4(a), 70% of identity requests were performed by sources having trust score higher or equal to 0.5. This percentage decreases to 61% in Figure 4(b), 56% in Figure 4(c) and approximately 50% in Figure 4(d).

In contrast, all sources being employed in the Sybil attack continue to present a discrepant behavior in regard to other sources. Despite the relative success attackers achieved when acting in collusion, these remain obtaining extremely low values of trust score (consequently, more complex computational puzzles). These results evidence the resilience and the efficacy of the proposed solution in face of Sybil attacks, even when these attacks occur in collusion. More importantly, the results show that the attacker needs to dedicate a large amount of resources to increase the gains obtained with the attack, both in terms of distributed sources (to outsmart the scheme of classification and aggregation of requests per source of origin), and in terms of computational capabilities (to solve the proposed puzzles).

V. FINAL CONSIDERATIONS

The use of computational puzzles is an alternative that has showed to be promising in tackling the occurrence of Sybil attacks in P2P networks. In spite of this, the lack of mechanisms to enable a proper treatment of situations in which there is a gap of computing power between legitimate users and attackers has hampered a more widespread and disseminated use of computational puzzles in the P2P realm. To address this limitation, in this paper we proposed the use of adaptive puzzles as a controlling factor to Sybil attacks.

The experiments carried out showed the efficacy of the proposed solution in decreasing the attackers’ capabilities of creating an indiscriminate number of counterfeit identities, whereas not compromising legitimate users, which were, in general, minimally penalized. When computing lower trust scores to sources having higher recurrence rates, (malicious) users associated to these sources had to cope with more complex computational puzzles. Conversely, users associated to presumably legitimate sources (and that made fewer use of the bootstrap service to request new identities) received less complex computational puzzles (given the higher values of trust scores determined for the great majority of these sources in the P2P network).

As prospective directions for future research, we intend to (i) investigate a mechanism to support the proper assignment of values to the various parameters of the proposed solution, taking into account communities having distinct characteristics; (ii) extend the evaluation of the proposed solution to capture the behavior of users in face of the delay associated to solving puzzles; and (iii) instantiate and evaluate the proposed solution as an extension of an existing P2P file sharing system, being BitTorrent a strong candidate.

ACKNOWLEDGEMENTS

The authors thank Prof. Nazareno Andrade (Federal University of Campina Grande, Brazil), for providing the historical traces of identity requests from Bitsoup.org used in the experimental evaluation presented in this paper.

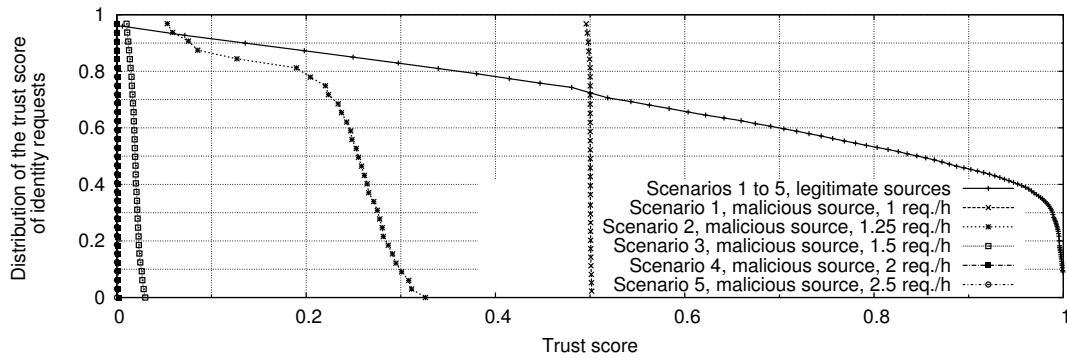


Fig. 3. Resilience of the proposed solution to Sybil attacks launched from a single malicious source, considering distinct recurrence rates

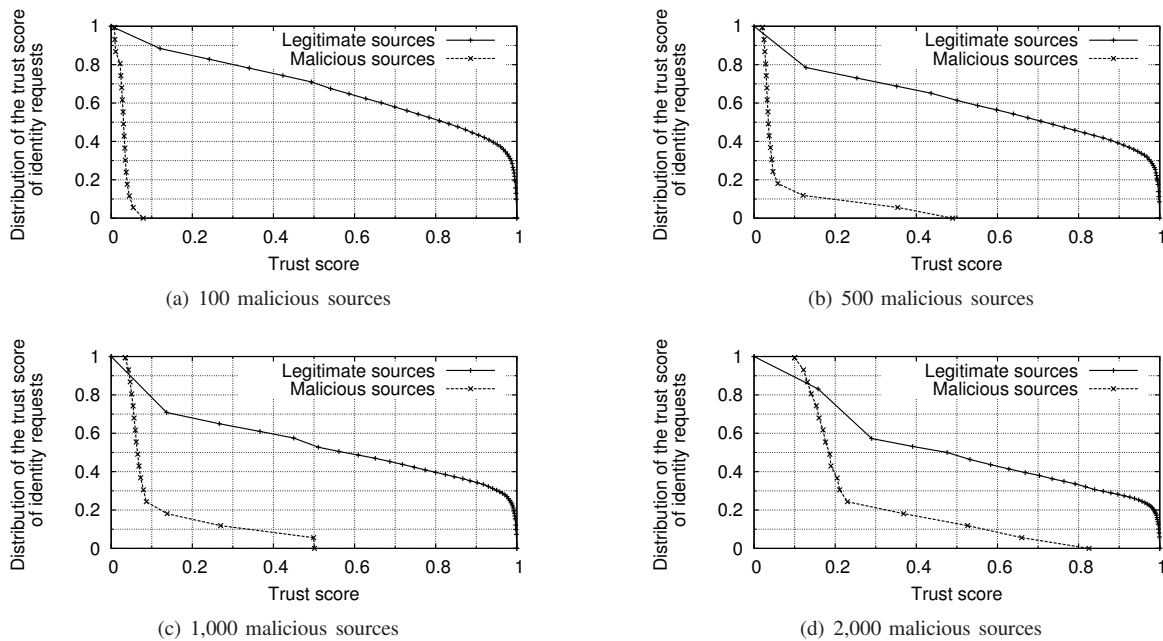


Fig. 4. Resilience of the proposed solution to Sybil attacks departing from multiple sources, considering a same recurrence rate

REFERENCES

- [1] J. R. Douceur, "The sybil attack," in *1st International Workshop on Peer-to-Peer Systems (IPTPS 2002)*, 2002, pp. 251–260.
- [2] M. P. Barcellos, D. Bauermann, H. Sant'anna, M. Lehmann, and R. Mansilha, "Protecting bittorrent: Design and evaluation of effective countermeasures against dos attacks," in *SRDS '08: Proceedings of the 2008 Symposium on Reliable Distributed Systems*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 73–82.
- [3] A. Singh, T.-W. Ngan, P. Druschel, and D. S. Wallach, "Eclipse attacks on overlay networks: Threats and defenses," in *25th Conference on Computer Communications (INFOCOM 2006)*, Barcelona, Catalunya, Spain, 2006, pp. 1–12.
- [4] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica, "Free-riding and whitewashing in peer-to-peer systems," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 1010–1019, 2006.
- [5] F. R. Santos, W. L. da Costa Cordeiro, L. P. Gaspary, and M. P. Barcellos, "Choking polluters in bittorrent file sharing communities," in *12th IFIP/IEEE Network Operations and Management Symposium (NOMS 2010)*, April 2010, pp. 1–8.
- [6] N. Borisov, "Computational puzzles as sybil defenses," in *6th IEEE International Conference on Peer-to-Peer Computing (P2P 2006)*, September 2006, pp. 171–176.
- [7] H. Rowaihy, W. Enck, P. McDaniel, and T. La Porta, "Limiting sybil attacks in structured p2p networks," in *26th IEEE International Conference on Computer Communications (INFOCOM 2007)*, Anchorage, Alaska, USA, May 2007, pp. 2596–2600.
- [8] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: a decentralized network coordinate system," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 4, pp. 15–26, October 2004.
- [9] M. Sherr, M. Blaze, and B. T. Loo, "Veracity: Practical Secure Network Coordinates via Vote-based Agreements," in *USENIX Annual Technical Conference (USENIX '09)*, June 2009.
- [10] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao, "Sybillimit: A near-optimal social network defense against sybil attacks," in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2008, pp. 3–17.
- [11] G. Danezis, C. Lesniewski-Laas, F. M. Kaashoek, and R. Anderson, "Sybil-resistant dht routing," 2005, pp. 305–318.
- [12] R. Morselli, B. Bhattacharjee, J. Katz, and M. A. Marsh, "Keychains: A decentralized public-key infrastructure," 2006.
- [13] K. Aberer, A. Datta, and M. Hauswirth, "A decentralized public key infrastructure for customer-to customer e-commerce," in *International Journal of Business Process Integration and Management*, 2005, pp. 26–33.
- [14] Bitsoup.org, "Bitsoup.org – the number one site for your torrent appetite," 2010.